GAME THEORY FOR APPLIED ECONOMISTS

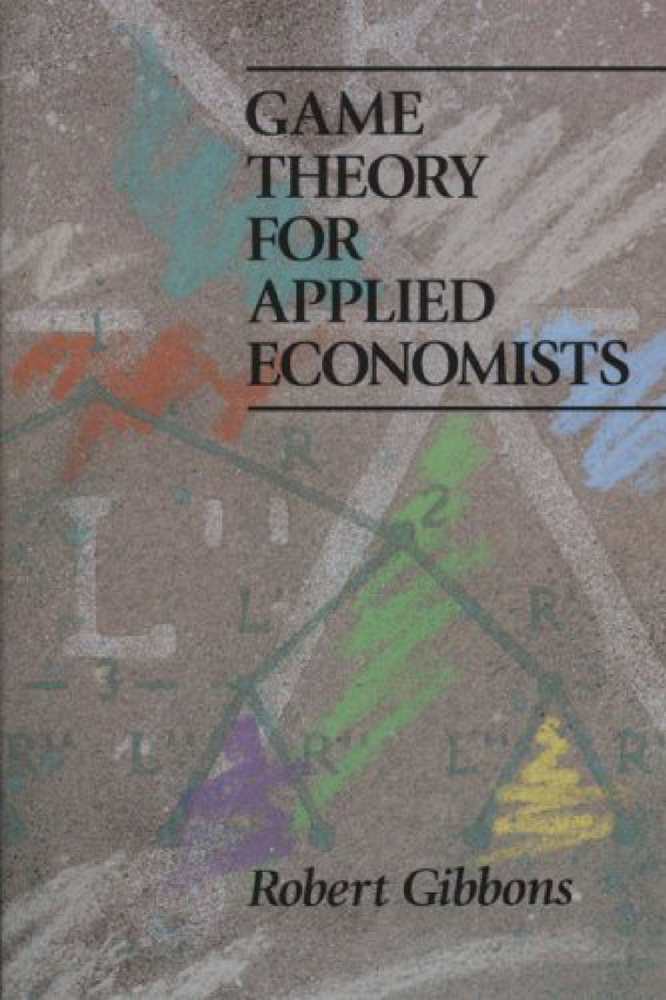

Robert Gibbons

Game Theory for Applied Economists

Robert Gibbons

Princeton University Press
Princeton, New Jersey

Contents

1 **Static Games of Complete Information** 1
  1.1 Basic Theory: Normal-Form Games and Nash Equilibrium 2
    1.1.A Normal-Form Representation of Games 2
    1.1.B Iterated Elimination of Strictly Dominated Strategies 4
    1.1.C Motivation and Definition of Nash Equilibrium 8
  1.2 Applications 14
    1.2.A Cournot Model of Duopoly 14
    1.2.B Bertrand Model of Duopoly 21
    1.2.C Final-Offer Arbitration 22
    1.2.D The Problem of the Commons 27
  1.3 Advanced Theory: Mixed Strategies and Existence of Equilibrium 29
    1.3.A Mixed Strategies 29
    1.3.B Existence of Nash Equilibrium 33
  1.4 Further Reading 48
  1.5 Problems 48
  1.6 References 51

2 **Dynamic Games of Complete Information** 55
  2.1 Dynamic Games of Complete and Perfect Information 57
    2.1.A Theory: Backwards Induction 57
    2.1.B Stackelberg Model of Duopoly 61
    2.1.C Wages and Employment in a Unionized Firm 64
    2.1.D Sequential Bargaining 68
  2.2 Two-Stage Games of Complete but Imperfect Information 71

2.2.A Theory: Subgame Perfection 71
2.2.B Bank Runs 73
2.2.C Tariffs and Imperfect International Competition . 75
2.2.D Tournaments 79
2.3 Repeated Games 82
2.3.A Theory: Two-Stage Repeated Games 82
2.3.B Theory: Infinitely Repeated Games 88
2.3.C Collusion between Cournot Duopolists 102
2.3.D Efficiency Wages 107
2.3.E Time-Consistent Monetary Policy . . . 112
2.4 Dynamic Games of Complete but Imperfect Information 115
2.4.A Extensive-Form Representation of Games . 115
2.4.B Subgame-Perfect Nash Equilibrium 122
2.5 Further Reading 129
2.6 Problems 130
2.7 References 138

3 Static Games of Incomplete Information 143
3.1 Theory: Static Bayesian Games and Bayesian Nash Equilibrium 144
3.1.A An Example: Cournot Competition under Asymmetric Information 144
3.1.B Normal-Form Representation of Static Bayesian Games 146
3.1.C Definition of Bayesian Nash Equilibrium . . 149
3.2 Applications 152
3.2.A Mixed Strategies Revisited 152
3.2.B An Auction 155
3.2.C A Double Auction 158
3.3 The Revelation Principle 164
3.4 Further Reading 168
3.5 Problems 169
3.6 References 172

4 Dynamic Games of Incomplete Information 173
4.1 Introduction to Perfect Bayesian Equilibrium 175
4.2 Signaling Games 183
4.2.A Perfect Bayesian Equilibrium in Signaling Games 183

4.2.B Job-Market Signaling 190
4.2.C Corporate Investment and Capital Structure . 205
4.2.D Monetary Policy 208
4.3 Other Applications of Perfect Bayesian Equilibrium 210
4.3.A Cheap-Talk Games 210
4.3.B Sequential Bargaining under Asymmetric Information 218
4.3.C Reputation in the Finitely Repeated Prisoners' Dilemma 224
4.4 Refinements of Perfect Bayesian Equilibrium 233
4.5 Further Reading 244
4.6 Problems 245
4.7 References 253

Index **257**

Preface

Game theory is the study of multiperson decision problems. Such problems arise frequently in economics. As is widely appreciated, for example, oligopolies present multiperson problems — each firm must consider what the others will do. But many other applications of game theory arise in fields of economics other than industrial organization. At the micro level, models of trading processes (such as bargaining and auction models) involve game theory. At an intermediate level of aggregation, labor and financial economics include game-theoretic models of the behavior of a firm in its input markets (rather than its output market, as in an oligopoly). There also are multiperson problems within a firm: many workers may vie for one promotion; several divisions may compete for the corporation's investment capital. Finally, at a high level of aggregation, international economics includes models in which countries compete (or collude) in choosing tariffs and other trade policies, and macroeconomics includes models in which the monetary authority and wage or price setters interact strategically to determine the effects of monetary policy.

This book is designed to introduce game theory to those who will later construct (or at least consume) game-theoretic models in applied fields within economics. The exposition emphasizes the economic applications of the theory at least as much as the pure theory itself, for three reasons. First, the applications help teach the theory; formal arguments about abstract games also appear but play a lesser role. Second, the applications illustrate the process of model building — the process of translating an informal description of a multiperson decision situation into a formal, game-theoretic problem to be analyzed. Third, the variety of applications shows that similar issues arise in different areas of economics, and that the same game-theoretic tools can be applied in each setting. In order to emphasize the broad potential scope of the theory, conventional applications from industrial organization largely have been replaced by applications from labor, macro, and other applied fields in economics.[1]

We will discuss four classes of games: static games of complete information, dynamic games of complete information, static games of incomplete information, and dynamic games of incomplete information. (A game has incomplete information if one player does not know another player's payoff, such as in an auction when one bidder does not know how much another bidder is willing to pay for the good being sold.) Corresponding to these four classes of games will be four notions of equilibrium in games: Nash equilibrium, subgame-perfect Nash equilibrium, Bayesian Nash equilibrium, and perfect Bayesian equilibrium.

Two (related) ways to organize one's thinking about these equilibrium concepts are as follows. First, one could construct sequences of equilibrium concepts of increasing strength, where stronger (i.e., more restrictive) concepts are attempts to eliminate implausible equilibria allowed by weaker notions of equilibrium. We will see, for example, that subgame-perfect Nash equilibrium is stronger than Nash equilibrium and that perfect Bayesian equilibrium in turn is stronger than subgame-perfect Nash equilibrium. Second, one could say that the equilibrium concept of interest is always perfect Bayesian equilibrium (or perhaps an even stronger equilibrium concept), but that it is equivalent to Nash equilibrium in static games of complete information, equivalent to subgame-perfection in dynamic games of complete (and perfect) information, and equivalent to Bayesian Nash equilibrium in static games of incomplete information.

The book can be used in two ways. For first-year graduate students in economics, many of the applications will already be familiar, so the game theory can be covered in a half-semester course, leaving many of the applications to be studied outside of class. For undergraduates, a full-semester course can present the theory a bit more slowly, as well as cover virtually all the applications in class. The main mathematical prerequisite is single-variable calculus; the rudiments of probability and analysis are introduced as needed.

---

[1] A good source for applications of game theory in industrial organization is Tirole's *The Theory of Industrial Organization* (MIT Press, 1988).

I learned game theory from David Kreps, John Roberts, and Bob Wilson in graduate school, and from Adam Brandenburger, Drew Fudenberg, and Jean Tirole afterward. I owe the theoretical perspective in this book to them. The focus on applications and other aspects of the pedagogical style, however, are largely due to the students in the MIT Economics Department from 1985 to 1990, who inspired and rewarded the courses that led to this book. I am very grateful for the insights and encouragement all these friends have provided, as well as for the many helpful comments on the manuscript I received from Joe Farrell, Milt Harris, George Mailath, Matthew Rabin, Andy Weiss, and several anonymous reviewers. Finally, I am glad to acknowledge the advice and encouragement of Jack Repcheck of Princeton University Press and financial support from an Olin Fellowship in Economics at the National Bureau of Economic Research.

Game Theory for Applied Economists

Chapter 1

Static Games of Complete Information

In this chapter we consider games of the following simple form: first the players simultaneously choose actions; then the players receive payoffs that depend on the combination of actions just chosen. Within the class of such static (or simultaneous-move) games, we restrict attention to games of *complete information*. That is, each player's payoff function (the function that determines the player's payoff from the combination of actions chosen by the players) is common knowledge among all the players. We consider dynamic (or sequential-move) games in Chapters 2 and 4, and games of incomplete information (games in which some player is uncertain about another player's payoff function—as in an auction where each bidder's willingness to pay for the good being sold is unknown to the other bidders) in Chapters 3 and 4.

In Section 1.1 we take a first pass at the two basic issues in game theory: how to describe a game and how to solve the resulting game-theoretic problem. We develop the tools we will use in analyzing static games of complete information, and also the foundations of the theory we will use to analyze richer games in later chapters. We define the *normal-form representation* of a game and the notion of a *strictly dominated strategy*. We show that some games can be solved by applying the idea that rational players do not play strictly dominated strategies, but also that in other games this approach produces a very imprecise prediction about the play of the game (sometimes as imprecise as "anything could happen"). We then motivate and define *Nash equilibrium*—a solution concept that produces much tighter predictions in a very broad class of games.

In Section 1.2 we analyze four applications, using the tools developed in the previous section: Cournot's (1838) model of imperfect competition, Bertrand's (1883) model of imperfect competition, Farber's (1980) model of final-offer arbitration, and the problem of the commons (discussed by Hume [1739] and others). In each application we first translate an informal statement of the problem into a normal-form representation of the game and then solve for the game's Nash equilibrium. (Each of these applications has a unique Nash equilibrium, but we discuss examples in which this is not true.)

In Section 1.3 we return to theory. We first define the notion of a *mixed strategy*, which we will interpret in terms of one player's uncertainty about what another player will do. We then state and discuss Nash's (1950) Theorem, which guarantees that a Nash equilibrium (possibly involving mixed strategies) exists in a broad class of games. Since we present first basic theory in Section 1.1, then applications in Section 1.2, and finally more theory in Section 1.3, it should be apparent that mastering the additional theory in Section 1.3 is not a prerequisite for understanding the applications in Section 1.2. On the other hand, the ideas of a mixed strategy and the existence of equilibrium do appear (occasionally) in later chapters.

This and each subsequent chapter concludes with problems, suggestions for further reading, and references.

1.1 Basic Theory: Normal-Form Games and Nash Equilibrium

1.1.A Normal-Form Representation of Games

In the normal-form representation of a game, each player simultaneously chooses a strategy, and the combination of strategies chosen by the players determines a payoff for each player. We illustrate the normal-form representation with a classic example — *The Prisoners' Dilemma.* Two suspects are arrested and charged with a crime. The police lack sufficient evidence to convict the suspects, unless at least one confesses. The police hold the suspects in separate cells and explain the consequences that will follow from the actions they could take. If neither confesses then both will be convicted of a minor offense and sentenced to one month in jail. If both confess then both will be sentenced to jail for six months. Finally, if one confesses but the other does not, then the confessor will be released immediately but the other will be sentenced to nine months in jail—six for the crime and a further three for obstructing justice.

The prisoners' problem can be represented in the accompanying bi-matrix. (Like a matrix, a bi-matrix can have an arbitrary number or rows and columns; "bi" refers to the fact that, in a two-player game, there are two numbers in each cell—the payoffs to the two players.)

		Prisoner 2	
		Mum	Fink
Prisoner 1	Mum	$-1, -1$	$-9, \ 0$
	Fink	$0, -9$	$-6, -6$

The Prisoners' Dilemma

In this game, each player has two strategies available: confess (or fink) and not confess (or be mum). The payoffs to the two players when a particular pair of strategies is chosen are given in the appropriate cell of the bi-matrix. By convention, the payoff to the so-called row player (here, Prisoner 1) is the first payoff given, followed by the payoff to the column player (here, Prisoner 2). Thus, if Prisoner 1 chooses Mum and Prisoner 2 chooses Fink, for example, then Prisoner 1 receives the payoff -9 (representing nine months in jail) and Prisoner 2 receives the payoff 0 (representing immediate release).

We now turn to the general case. The *normal-form representation* of a game specifies: (1) the players in the game, (2) the strategies available to each player, and (3) the payoff received by each player for each combination of strategies that could be chosen by the players. We will often discuss an n-player game in which the players are numbered from 1 to n and an arbitrary player is called player i. Let S_i denote the set of strategies available to player i (called i's *strategy space*), and let s_i denote an arbitrary member of this set. (We will occasionally write $s_i \in S_i$ to indicate that the strategy s_i is a member of the set of strategies S_i.) Let $(s_1, \ldots, s_n)$ denote a combination of strategies, one for each player, and let u_i denote player i's *payoff function*: $u_i(s_1, \ldots, s_n)$ is the payoff to player i if the players choose the strategies $(s_1, \ldots, s_n)$. Collecting all of this information together, we have:

Definition *The* ***normal-form representation*** *of an n-player game specifies the players' strategy spaces* $S_1, \ldots, S_n$ *and their payoff functions* $u_1, \ldots, u_n$. *We denote this game by* $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$.

Although we stated that in a normal-form game the players choose their strategies simultaneously, this does not imply that the parties necessarily *act* simultaneously: it suffices that each choose his or her action without knowledge of the others' choices, as would be the case here if the prisoners reached decisions at arbitrary times while in their separate cells. Furthermore, although in this chapter we use normal-form games to represent only static games in which the players all move without knowing the other players' choices, we will see in Chapter 2 that normal-form representations can be given for sequential-move games, but also that an alternative—the *extensive-form* representation of the game—is often a more convenient framework for analyzing dynamic issues.

1.1.B Iterated Elimination of Strictly Dominated Strategies

Having described one way to represent a game, we now take a first pass at describing how to solve a game-theoretic problem. We start with the Prisoners' Dilemma because it is easy to solve, using only the idea that a rational player will not play a strictly dominated strategy.

In the Prisoners' Dilemma, if one suspect is going to play Fink, then the other would prefer to play Fink and so be in jail for six months rather than play Mum and so be in jail for nine months. Similarly, if one suspect is going to play Mum, then the other would prefer to play Fink and so be released immediately rather than play Mum and so be in jail for one month. Thus, for prisoner i, playing Mum is dominated by playing Fink—for each strategy that prisoner j could choose, the payoff to prisoner i from playing Mum is less than the payoff to i from playing Fink. (The same would be true in any bi-matrix in which the payoffs 0, -1, -6, and -9 above were replaced with payoffs T, R, P, and S, respectively, provided that $T > R > P > S$ so as to capture the ideas of temptation, reward, punishment, and sucker payoffs.) More generally:

Definition *In the normal-form game* $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$, *let* s_i' *and* s_i'' *be feasible strategies for player* i *(i.e.,* s_i' *and* s_i'' *are members of* S_i*). Strategy* s_i' *is* ***strictly dominated*** *by strategy* s_i'' *if for each feasible combination of the other players' strategies,* i*'s payoff from playing* s_i' *is strictly less than* i*'s payoff from playing* s_i''*:*

$$u_i(s_1, \ldots, s_{i-1}, s_i', s_{i+1}, \ldots, s_n) < u_i(s_1, \ldots, s_{i-1}, s_i'', s_{i+1}, \ldots, s_n) \quad \text{(DS)}$$

for each $(s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ *that can be constructed from the other players' strategy spaces* $S_1, \ldots, S_{i-1}, S_{i+1}, \ldots, S_n$.

Rational players do not play strictly dominated strategies, because there is no belief that a player could hold (about the strategies the other players will choose) such that it would be optimal to play such a strategy.[1] Thus, in the Prisoners' Dilemma, a rational player will choose Fink, so (Fink, Fink) will be the outcome reached by two rational players, even though (Fink, Fink) results in worse payoffs for both players than would (Mum, Mum). Because the Prisoners' Dilemma has many applications (including the arms race and the free-rider problem in the provision of public goods), we will return to variants of the game in Chapters 2 and 4. For now, we focus instead on whether the idea that rational players do not play strictly dominated strategies can lead to the solution of other games.

Consider the abstract game in Figure 1.1.1.[2] Player 1 has two strategies and player 2 has three: $S_1 = \{\text{Up}, \text{Down}\}$ and $S_2 = \{\text{Left}, \text{Middle}, \text{Right}\}$. For player 1, neither Up nor Down is strictly

---

[1] A complementary question is also of interest: if there is no belief that player i could hold (about the strategies the other players will choose) such that it would be optimal to play the strategy s_i, can we conclude that there must be another strategy that strictly dominates s_i? The answer is "yes," *provided* that we adopt appropriate definitions of "belief" and "another strategy," both of which involve the idea of mixed strategies to be introduced in Section 1.3.A.

[2] Most of this book considers economic applications rather than abstract examples, both because the applications are of interest in their own right and because, for many readers, the applications are often a useful way to explain the underlying theory. When introducing some of the basic theoretical ideas, however, we will sometimes resort to abstract examples that have no natural economic interpretation.

		Player 2		
		Left	Middle	Right
Player 1	Up	1,0	1,2	0,1
	Down	0,3	0,1	2,0

Figure 1.1.1.

dominated: Up is better than Down if 2 plays Left (because $1 > 0$), but Down is better than Up if 2 plays Right (because $2 > 0$). For player 2, however, Right is strictly dominated by Middle (because $2 > 1$ and $1 > 0$), so a rational player 2 will not play Right. Thus, if player 1 knows that player 2 is rational then player 1 can eliminate Right from player 2's strategy space. That is, if player 1 knows that player 2 is rational then player 1 can play the game in Figure 1.1.1 *as if* it were the game in Figure 1.1.2.

		Player 2	
		Left	Middle
Player 1	Up	1,0	1,2
	Down	0,3	0,1

Figure 1.1.2.

In Figure 1.1.2, Down is now strictly dominated by Up for player 1, so if player 1 is rational (and player 1 knows that player 2 is rational, so that the game in Figure 1.1.2 applies) then player 1 will not play Down. Thus, if player 2 knows that player 1 is rational, *and* player 2 knows that player 1 knows that player 2 is rational (so that player 2 knows that Figure 1.1.2 applies), then player 2 can eliminate Down from player 1's strategy space, leaving the game in Figure 1.1.3. But now Left is strictly dominated by Middle for player 2, leaving (Up, Middle) as the outcome of the game.

This process is called *iterated elimination of strictly dominated strategies.* Although it is based on the appealing idea that rational players do not play strictly dominated strategies, the process has two drawbacks. First, each step requires a further assumption

		Player 2	
		Left	Middle
Player 1	Up	1,0	1,2

Figure 1.1.3.

about what the players know about each other's rationality. If we want to be able to apply the process for an arbitrary number of steps, we need to assume that it is *common knowledge* that the players are rational. That is, we need to assume not only that all the players are rational, but also that all the players know that all the players are rational, and that all the players know that all the players know that all the players are rational, and so on, *ad infinitum.* (See Aumann [1976] for the formal definition of common knowledge.)

The second drawback of iterated elimination of strictly dominated strategies is that the process often produces a very imprecise prediction about the play of the game. Consider the game in Figure 1.1.4, for example. In this game there are no strictly dominated strategies to be eliminated. (Since we have not motivated this game in the slightest, it may appear arbitrary, or even pathological. See the case of three or more firms in the Cournot model in Section 1.2.A for an economic application in the same spirit.) Since all the strategies in the game survive iterated elimination of strictly dominated strategies, the process produces no prediction whatsoever about the play of the game.

	L	C	R
T	0,4	4,0	5,3
M	4,0	0,4	5,3
B	3,5	3,5	6,6

Figure 1.1.4.

We turn next to Nash equilibrium—a solution concept that produces much tighter predictions in a very broad class of games. We show that Nash equilibrium is a stronger solution concept than iterated elimination of strictly dominated strategies, in the sense that the players' strategies in a Nash equilibrium always survive iterated elimination of strictly dominated strategies, but the converse is not true. In subsequent chapters we will argue that in richer games even Nash equilibrium produces too imprecise a prediction about the play of the game, so we will define still-stronger notions of equilibrium that are better suited for these richer games.

1.1.C Motivation and Definition of Nash Equilibrium

One way to motivate the definition of Nash equilibrium is to argue that if game theory is to provide a unique solution to a game-theoretic problem then the solution must be a Nash equilibrium, in the following sense. Suppose that game theory makes a unique prediction about the strategy each player will choose. In order for this prediction to be correct, it is necessary that each player be willing to choose the strategy predicted by the theory. Thus, each player's predicted strategy must be that player's best response to the predicted strategies of the other players. Such a prediction could be called *strategically stable* or *self-enforcing*, because no single player wants to deviate from his or her predicted strategy. We will call such a prediction a Nash equilibrium:

Definition *In the n-player normal-form game* $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$, *the strategies* $(s_1^*, \ldots, s_n^*)$ *are a* ***Nash equilibrium*** *if, for each player* i, s_i^* *is (at least tied for) player i's best response to the strategies specified for the* $n-1$ *other players,* $(s_1^*, \ldots . s_{i-1}^*, s_{i+1}^*, \ldots, s_n^*)$:

$$\begin{aligned} & u_i(s_1^*, \ldots, s_{i-1}^*, s_i^*, s_{i+1}^*, \ldots, s_n^*) \\ & \qquad \geq u_i(s_1^*, \ldots, s_{i-1}^*, s_i, s_{i+1}^*, \ldots, s_n^*) \end{aligned} \tag{NE}$$

for every feasible strategy s_i *in* S_i; *that is,* s_i^* *solves*

$$\max_{s_i \in S_i} u_i(s_1^*, \ldots, s_{i-1}^*, s_i, s_{i+1}^*, \ldots, s_n^*).$$

To relate this definition to its motivation, suppose game theory offers the strategies $(s_1', \ldots, s_n')$ as the solution to the normal-form game $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$. Saying that $(s_1', \ldots, s_n')$ is *not* a Nash equilibrium of G is equivalent to saying that there exists some player i such that s'_i is *not* a best response to $(s'_1, \ldots, s'_{i-1}, s'_{i+1}, \ldots, s'_n)$. That is, there exists some s''_i in S_i such that

$$u_i(s'_1, \ldots, s'_{i-1}, s'_i, s'_{i+1}, \ldots, s'_n) < u_i(s'_1, \ldots, s'_{i-1}, s''_i, s'_{i+1}, \ldots, s'_n).$$

Thus, if the theory offers the strategies $(s'_1, \ldots, s'_n)$ as the solution but these strategies are not a Nash equilibrium, then at least one player will have an incentive to deviate from the theory's prediction, so the theory will be falsified by the actual play of the game. A closely related motivation for Nash equilibrium involves the idea of convention: if a convention is to develop about how to play a given game then the strategies prescribed by the convention must be a Nash equilibrium, else at least one player will not abide by the convention.

To be more concrete, we now solve a few examples. Consider the three normal-form games already described—the Prisoners' Dilemma and Figures 1.1.1 and 1.1.4. A brute-force approach to finding a game's Nash equilibria is simply to check whether each possible combination of strategies satisfies condition (NE) in the definition.[3] In a two-player game, this approach begins as follows: for each player, and for each feasible strategy for that player, determine the other player's best response to that strategy. Figure 1.1.5 does this for the game in Figure 1.1.4 by underlining the payoff to player j's best response to each of player i's feasible strategies. If the column player were to play L, for instance, then the row player's best response would be M, since 4 exceeds 3 and 0, so the row player's payoff of 4 in the (M, L) cell of the bi-matrix is underlined.

A pair of strategies satisfies condition (NE) if each player's strategy is a best response to the other's—that is, if both payoffs are underlined in the corresponding cell of the bi-matrix. Thus, (B, R) is the only strategy pair that satisfies (NE); likewise for (Fink, Fink) in the Prisoners' Dilemma and (Up, Middle) in

---

[3] In Section 1.3.A we will distinguish between pure and mixed strategies. We will then see that the definition given here describes *pure-strategy* Nash equilibria, but that there can also be *mixed-strategy* Nash equilibria. Unless explicitly noted otherwise, all references to Nash equilibria in this section are to pure-strategy Nash equilibria.

	L	C	R
T	$0,\underline{4}$	$\underline{4},0$	$5,3$
M	$\underline{4},0$	$0,\underline{4}$	$5,3$
B	$3,5$	$3,5$	$\underline{6},\underline{6}$

Figure 1.1.5.

Figure 1.1.1. These strategy pairs are the unique Nash equilibria of these games.[4]

We next address the relation between Nash equilibrium and iterated elimination of strictly dominated strategies. Recall that the Nash equilibrium strategies in the Prisoners' Dilemma and Figure 1.1.1—(Fink, Fink) and (Up, Middle), respectively—are the only strategies that survive iterated elimination of strictly dominated strategies. This result can be generalized: if iterated elimination of strictly dominated strategies eliminates all but the strategies $(s_1^*,\ldots,s_n^*)$, then these strategies are the unique Nash equilibrium of the game. (See Appendix 1.1.C for a proof of this claim.) Since iterated elimination of strictly dominated strategies frequently does *not* eliminate all but a single combination of strategies, however, it is of more interest that Nash equilibrium is a stronger solution concept than iterated elimination of strictly dominated strategies, in the following sense. If the strategies $(s_1^*,\ldots,s_n^*)$ are a Nash equilibrium then they survive iterated elimination of strictly dominated strategies (again, see the Appendix for a proof), but there can be strategies that survive iterated elimination of strictly dominated strategies but are not part of any Nash equilibrium. To see the latter, recall that in Figure 1.1.4 Nash equilibrium gives the unique prediction (B, R), whereas iterated elimination of strictly dominated strategies gives the maximally imprecise prediction: no strategies are eliminated; anything could happen.

Having shown that Nash equilibrium is a stronger solution concept than iterated elimination of strictly dominated strategies, we must now ask whether Nash equilibrium is too strong a solution concept. That is, can we be sure that a Nash equilibrium exists? Nash (1950) showed that in any finite game (i.e., a game in which the number of players n and the strategy sets $S_1, \ldots, S_n$ are all finite) there exists at least one Nash equilibrium. (This equilibrium may involve mixed strategies, which we will discuss in Section 1.3.A; see Section 1.3.B for a precise statement of Nash's Theorem.) Cournot (1838) proposed the same notion of equilibrium in the context of a particular model of duopoly and demonstrated (by construction) that an equilibrium exists in that model; see Section 1.2.A. In every application analyzed in this book, we will follow Cournot's lead: we will demonstrate that a Nash (or stronger) equilibrium exists by constructing one. In some of the theoretical sections, however, we will rely on Nash's Theorem (or its analog for stronger equilibrium concepts) and simply assert that an equilibrium exists.

---

[4] This statement is correct even if we do not restrict attention to pure-strategy Nash equilibrium, because no mixed-strategy Nash equilibria exist in these three games. See Problem 1.10.

We conclude this section with another classic example—*The Battle of the Sexes*. This example shows that a game can have multiple Nash equilibria, and also will be useful in the discussions of mixed strategies in Sections 1.3.B and 3.2.A. In the traditional exposition of the game (which, it will be clear, dates from the 1950s), a man and a woman are trying to decide on an evening's entertainment; we analyze a gender-neutral version of the game. While at separate workplaces, Pat and Chris must choose to attend either the opera or a prize fight. Both players would rather spend the evening together than apart, but Pat would rather they be together at the prize fight while Chris would rather they be together at the opera, as represented in the accompanying bi-matrix.

		Pat	
		Opera	Fight
Chris	Opera	2,1	0,0
	Fight	0,0	1,2

The Battle of the Sexes

Both (Opera, Opera) and (Fight, Fight) are Nash equilibria.

We argued above that if game theory is to provide a unique solution to a game then the solution must be a Nash equilibrium. This argument ignores the possibility of games in which game theory does not provide a unique solution. We also argued that if a convention is to develop about how to play a given game, then the strategies prescribed by the convention must be a Nash equilibrium, but this argument similarly ignores the possibility of games for which a convention will not develop. In some games with multiple Nash equilibria one equilibrium stands out as the compelling solution to the game. (Much of the theory in later chapters is an effort to identify such a compelling equilibrium in different classes of games.) Thus, the existence of multiple Nash equilibria is not a problem in and of itself. In the Battle of the Sexes, however, (Opera, Opera) and (Fight, Fight) seem equally compelling, which suggests that there may be games for which game theory does not provide a unique solution and no convention will develop.[5] In such games, Nash equilibrium loses much of its appeal as a prediction of play.

Appendix 1.1.C

This appendix contains proofs of the following two Propositions, which were stated informally in Section 1.1.C. Skipping these proofs will not substantially hamper one's understanding of later material. For readers not accustomed to manipulating formal definitions and constructing proofs, however, mastering these proofs will be a valuable exercise.

Proposition A *In the n-player normal-form game* $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$, *if iterated elimination of strictly dominated strategies eliminates all but the strategies* $(s_1^*, \ldots, s_n^*)$, *then these strategies are the unique Nash equilibrium of the game.*

Proposition B *In the n-player normal-form game* $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$, *if the strategies* $(s_1^*, \ldots, s_n^*)$ *are a Nash equilibrium, then they survive iterated elimination of strictly dominated strategies.*

---

[5] In Section 1.3.B we describe a third Nash equilibrium of the Battle of the Sexes (involving mixed strategies). Unlike (Opera, Opera) and (Fight, Fight), this third equilibrium has symmetric payoffs, as one might expect from the unique solution to a symmetric game; on the other hand, the third equilibrium is also inefficient, which may work against its development as a convention. Whatever one's judgment about the Nash equilibria in the Battle of the Sexes, however, the broader point remains: there may be games in which game theory does not provide a unique solution and no convention will develop.

Since Proposition B is simpler to prove, we begin with it, to warm up. The argument is by contradiction. That is, we will assume that one of the strategies in a Nash equilibrium is eliminated by iterated elimination of strictly dominated strategies, and then we will show that a contradiction would result if this assumption were true, thereby proving that the assumption must be false.

Suppose that the strategies $(s_1^*, \ldots, s_n^*)$ are a Nash equilibrium of the normal-form game $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$, but suppose also that (perhaps after some strategies other than $(s_1^*, \ldots, s_n^*)$ have been eliminated) s_i^* is the first of the strategies $(s_1^*, \ldots, s_n^*)$ to be eliminated for being strictly dominated. Then there must exist a strategy s_i'' that has not yet been eliminated from S_i that strictly dominates s_i^*. Adapting (DS), we have

$$\begin{aligned} &u_i(s_1, \ldots, s_{i-1}, s_i^*, s_{i+1}, \ldots, s_n) \\ &\qquad < u_i(s_1, \ldots, s_{i-1}, s_i'', s_{i+1}, \ldots, s_n) \end{aligned} \tag{1.1.1}$$

for each $(s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ that can be constructed from the strategies that have not yet been eliminated from the other players' strategy spaces. Since s_i^* is the first of the equilibrium strategies to be eliminated, the other players' equilibrium strategies have not yet been eliminated, so one of the implications of (1.1.1) is

$$\begin{aligned} &u_i(s_1^*, \ldots, s_{i-1}^*, s_i^*, s_{i+1}^*, \ldots, s_n^*) \\ &\qquad < u_i(s_1^*, \ldots, s_{i-1}^*. s_i'', s_{i+1}^*, \ldots, s_n^*). \end{aligned} \tag{1.1.2}$$

But (1.1.2) is contradicted by (NE): s_i^* must be a best response to $(s_1^*, \ldots, s_{i-1}^*, s_{i+1}^*, \ldots, s_n^*)$, so there cannot exist a strategy s_i'' that strictly dominates s_i^*. This contradiction completes the proof.

Having proved Proposition B, we have already proved part of Proposition A: all we need to show is that if iterated elimination of dominated strategies eliminates all but the strategies $(s_1^*, \ldots, s_n^*)$ then these strategies are a Nash equilibrium; by Proposition B, any other Nash equilibria would also have survived, so this equilibrium must be unique. We assume that G is finite.

The argument is again by contradiction. Suppose that iterated elimination of dominated strategies eliminates all but the strategies $(s_1^*, \ldots, s_n^*)$ but these strategies are not a Nash equilibrium. Then there must exist some player i and some feasible strategy s_i in S_i such that (NE) fails, but s_i must have been strictly dominated by some other strategy s_i' at some stage of the process. The formal statements of these two observations are: there exists s_i in S_i such that

$$u_i(s_1^*, \ldots, s_{i-1}^*, s_i^*, s_{i+1}^*, \ldots, s_n^*) \\ < u_i(s_1^*, \ldots, s_{i-1}^*, s_i, s_{i+1}^*, \ldots, s_n^*); \qquad (1.1.3)$$

and there exists s_i' in the set of player i's strategies remaining at some stage of the process such that

$$u_i(s_1, \ldots, s_{i-1}, s_i, s_{i+1}, \ldots, s_n) \\ < u_i(s_1, \ldots, s_{i-1}, s_i', s_{i+1}, \ldots, s_n) \qquad (1.1.4)$$

for each $(s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ that can be constructed from the strategies remaining in the other players' strategy spaces at that stage of the process. Since the other players' strategies $(s_1^*, \ldots, s_{i-1}^*, s_{i+1}^*, \ldots, s_n^*)$ are never eliminated, one of the implications of (1.1.4) is

$$u_i(s_1^*, \ldots, s_{i-1}^*, s_i, s_{i+1}^*, \ldots, s_n^*) \\ < u_i(s_1^*, \ldots, s_{i-1}^*, s_i', s_{i+1}^*, \ldots, s_n^*). \qquad (1.1.5)$$

If $s_i' = s_i^*$ (that is, if s_i^* is the strategy that strictly dominates s_i) then (1.1.5) contradicts (1.1.3), in which case the proof is complete. If $s_i' \neq s_i^*$ then some other strategy s_i'' must later strictly dominate s_i', since s_i' does not survive the process. Thus, inequalities analogous to (1.1.4) and (1.1.5) hold with s_i' and s_i'' replacing s_i and s_i', respectively. Once again, if $s_i'' = s_i^*$ then the proof is complete; otherwise, two more analogous inequalities can be constructed. Since s_i^* is the only strategy from S_i to survive the process, repeating this argument (in a finite game) eventually completes the proof.

1.2 Applications

1.2.A Cournot Model of Duopoly

As noted in the previous section, Cournot (1838) anticipated Nash's definition of equilibrium by over a century (but only in the context of a particular model of duopoly). Not surprisingly, Cournot's work is one of the classics of game theory; it is also one of the cornerstones of the theory of industrial organization. We consider a very simple version of Cournot's model here, and return to variations on the model in each subsequent chapter. In this section we use the model to illustrate: (a) the translation of an informal statement of a problem into a normal-form representation of a game; (b) the computations involved in solving for the game's Nash equilibrium; and (c) iterated elimination of strictly dominated strategies.

Let q_1 and q_2 denote the quantities (of a homogeneous product) produced by firms 1 and 2, respectively. Let $P(Q) = a - Q$ be the market-clearing price when the aggregate quantity on the market is $Q = q_1 + q_2$. (More precisely, $P(Q) = a - Q$ for $Q < a$, and $P(Q) = 0$ for $Q \geq a$.) Assume that the total cost to firm i of producing quantity q_i is $C_i(q_i) = cq_i$. That is, there are no fixed costs and the marginal cost is constant at c, where we assume $c < a$. Following Cournot, suppose that the firms choose their quantities simultaneously.[6]

In order to find the Nash equilibrium of the Cournot game, we first translate the problem into a normal-form game. Recall from the previous section that the normal-form representation of a game specifies: (1) the players in the game, (2) the strategies available to each player, and (3) the payoff received by each player for each combination of strategies that could be chosen by the players. There are of course two players in any duopoly game—the two firms. In the Cournot model, the strategies available to each firm are the different quantities it might produce. We will assume that output is continuously divisible. Naturally, negative outputs are not feasible. Thus, each firm's strategy space can be represented as $S_i = [0, \infty)$, the nonnegative real numbers, in which case a typical strategy s_i is a quantity choice, $q_i \geq 0$. One could argue that extremely large quantities are not feasible and so should not be included in a firm's strategy space. Because $P(Q) = 0$ for $Q \geq a$, however, neither firm will produce a quantity $q_i > a$.

It remains to specify the payoff to firm i as a function of the strategies chosen by it and by the other firm, and to define and solve for equilibrium. We assume that the firm's payoff is simply its profit. Thus, the payoff $u_i(s_i, s_j)$ in a general two-player game in normal form can be written here as[7]

---

[6] We discuss Bertrand's (1883) model, in which firms choose prices rather than quantities, in Section 1.2.B, and Stackelberg's (1934) model, in which firms choose quantities but one firm chooses before (and is observed by) the other, in Section 2.1.B. Finally, we discuss Friedman's (1971) model, in which the interaction described in Cournot's model occurs repeatedly over time, in Section 2.3.C.

$$\pi_i(q_i, q_j) = q_i[P(q_i + q_j) - c] = q_i[a - (q_i + q_j) - c].$$

Recall from the previous section that in a two-player game in normal form, the strategy pair (s_1^*, s_2^*) is a Nash equilibrium if, for each player i,

$$u_i(s_i^*, s_j^*) \geq u_i(s_i, s_j^*) \tag{NE}$$

for every feasible strategy s_i in S_i. Equivalently, for each player i, s_i^* must solve the optimization problem

$$\max_{s_i \in S_i} u_i(s_i, s_j^*).$$

In the Cournot duopoly model, the analogous statement is that the quantity pair (q_1^*, q_2^*) is a Nash equilibrium if, for each firm i, q_i^* solves

$$\max_{0 \leq q_i < \infty} \pi_i(q_i, q_j^*) = \max_{0 \leq q_i < \infty} q_i[a - (q_i + q_j^*) - c].$$

Assuming $q_j^* < a - c$ (as will be shown to be true), the first-order condition for firm i's optimization problem is both necessary and sufficient; it yields

$$q_i = \frac{1}{2}(a - q_j^* - c). \tag{1.2.1}$$

Thus, if the quantity pair (q_1^*, q_2^*) is to be a Nash equilibrium, the firms' quantity choices must satisfy

$$q_1^* = \frac{1}{2}(a - q_2^* - c)$$

and

$$q_2^* = \frac{1}{2}(a - q_1^* - c).$$

---

[7] Note that we have changed the notation slightly by writing $u_i(s_i, s_j)$ rather than $u_i(s_1, s_2)$. Both expressions represent the payoff to player i as a function of the strategies chosen by all the players. We will use these expressions (and their n-player analogs) interchangeably.

Solving this pair of equations yields

$$q_1^* = q_2^* = \frac{a-c}{3},$$

which is indeed less than $a - c$, as assumed.

The intuition behind this equilibrium is simple. Each firm would of course like to be a monopolist in this market, in which case it would choose q_i to maximize $\pi_i(q_i, 0)$—it would produce the monopoly quantity $q_m = (a-c)/2$ and earn the monopoly profit $\pi_i(q_m, 0) = (a-c)^2/4$. Given that there are two firms, aggregate profits for the duopoly would be maximized by setting the aggregate quantity $q_1 + q_2$ equal to the monopoly quantity q_m, as would occur if $q_i = q_m/2$ for each i, for example. The problem with this arrangement is that each firm has an incentive to deviate: because the monopoly quantity is low, the associated price $P(q_m)$ is high, and at this price each firm would like to increase its quantity, in spite of the fact that such an increase in production drives down the market-clearing price. (To see this formally, use (1.2.1) to check that $q_m/2$ is *not* firm 2's best response to the choice of $q_m/2$ by firm 1.) In the Cournot equilibrium, in contrast, the aggregate quantity is higher, so the associated price is lower, so the temptation to increase output is reduced—reduced by just enough that each firm is just deterred from increasing its output by the realization that the market-clearing price will fall. See Problem 1.4 for an analysis of how the presence of n oligopolists affects this equilibrium trade-off between the temptation to increase output and the reluctance to reduce the market-clearing price.

Rather than solving for the Nash equilibrium in the Cournot game algebraically, one could instead proceed graphically, as follows. Equation (1.2.1) gives firm i's best response to firm j's *equilibrium* strategy, q_j^*. Analogous reasoning leads to firm 2's best response to an *arbitrary* strategy by firm 1 and firm 1's best response to an arbitrary strategy by firm 2. Assuming that firm 1's strategy satisfies $q_1 < a - c$, firm 2's best response is

$$R_2(q_1) = \frac{1}{2}(a - q_1 - c);$$

likewise, if $q_2 < a - c$ then firm 1's best response is

$$R_1(q_2) = \frac{1}{2}(a - q_2 - c).$$

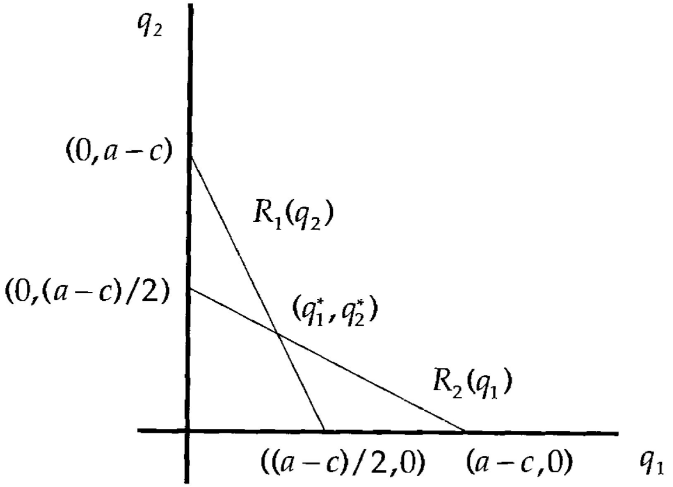

Figure 1.2.1.

As shown in Figure 1.2.1, these two best-response functions intersect only once, at the equilibrium quantity pair (q_1^*, q_2^*).

A third way to solve for this Nash equilibrium is to apply the process of iterated elimination of strictly dominated strategies. This process yields a unique solution—which, by Proposition A in Appendix 1.1.C, must be the Nash equilibrium (q_1^*, q_2^*). The complete process requires an infinite number of steps, each of which eliminates a fraction of the quantities remaining in each firm's strategy space; we discuss only the first two steps. First, the monopoly quantity $q_m = (a-c)/2$ strictly dominates any higher quantity. That is, for any $x > 0$, $\pi_i(q_m, q_j) > \pi_i(q_m + x, q_j)$ for all $q_j \geq 0$. To see this, note that if $Q = q_m + x + q_j < a$, then

$$\pi_i(q_m, q_j) = \frac{a-c}{2}\left[\frac{a-c}{2} - q_j\right]$$

and

$$\pi_i(q_m + x, q_j) = \left[\frac{a-c}{2} + x\right]\left[\frac{a-c}{2} - x - q_j\right] = \pi_i(q_m, q_j) - x(x + q_j),$$

and if $Q = q_m + x + q_j \geq a$, then $P(Q) = 0$, so producing a smaller quantity raises profit. Second, given that quantities exceeding q_m have been eliminated, the quantity $(a-c)/4$ strictly dominates any lower quantity. That is, for any x between zero and $(a-c)/4$, $\pi_i[(a-c)/4, q_j] > \pi_i[(a-c)/4 - x, q_j]$ for all q_j between zero and $(a-c)/2$. To see this, note that

$$\pi_i\left(\frac{a-c}{4}, q_j\right) = \frac{a-c}{4}\left[\frac{3(a-c)}{4} - q_j\right]$$

and

$$\begin{aligned}\pi_i\left(\frac{a-c}{4} - x, q_j\right) &= \left[\frac{a-c}{4} - x\right]\left[\frac{3(a-c)}{4} + x - q_j\right] \\ &= \pi_i(q_m, q_j) - x\left[\frac{a-c}{2} + x - q_j\right].\end{aligned}$$

After these two steps, the quantities remaining in each firm's strategy space are those in the interval between $(a-c)/4$ and $(a-c)/2$. Repeating these arguments leads to ever-smaller intervals of remaining quantities. In the limit, these intervals converge to the single point $q_i^* = (a-c)/3$.

Iterated elimination of strictly dominated strategies can also be described graphically, by using the observation (from footnote 1; see also the discussion in Section 1.3.A) that a strategy is strictly dominated if and only if there is no belief about the other players' choices for which the strategy is a best response. Since there are only two firms in this model, we can restate this observation as: a quantity q_i is strictly dominated if and only if there is no belief about q_j such that q_i is firm i's best response. We again discuss only the first two steps of the iterative process. First, it is never a best response for firm i to produce more than the monopoly quantity, $q_m = (a-c)/2$. To see this, consider firm 2's best-response function, for example: in Figure 1.2.1, $R_2(q_1)$ equals q_m when $q_1 = 0$ and declines as q_1 increases. Thus, for any $q_j \geq 0$, if firm i believes that firm j will choose q_j, then firm i's best response is less than or equal to q_m; there is no q_j such that firm i's best response exceeds q_m. Second, given this upper bound on firm j's quantity, we can derive a lower bound on firm i's best response: if $q_j \leq (a-c)/2$, then $R_i(q_j) \geq (a-c)/4$, as shown for firm 2's best response in Figure 1.2.2.[8]

---

[8] These two arguments are slightly incomplete because we have not analyzed

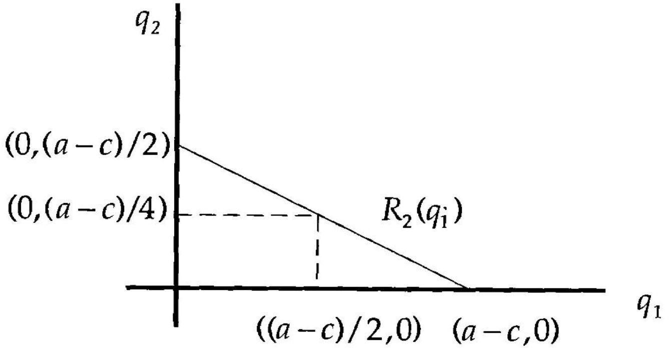

Figure 1.2.2.

As before, repeating these arguments leads to the single quantity $q_i^* = (a-c)/3$.

We conclude this section by changing the Cournot model so that iterated elimination of strictly dominated strategies does *not* yield a unique solution. To do this, we simply add one or more firms to the existing duopoly. We will see that the first of the two steps discussed in the duopoly case continues to hold, but that the process ends there. Thus, when there are more than two firms, iterated elimination of strictly dominated strategies yields only the imprecise prediction that each firm's quantity will not exceed the monopoly quantity (much as in Figure 1.1.4, where no strategies were eliminated by this process).

For concreteness, we consider the three-firm case. Let Q_{-i} denote the sum of the quantities chosen by the firms other than i, and let $\pi_i(q_i, Q_{-i}) = q_i(a - q_i - Q_{-i} - c)$ provided $q_i + Q_{-i} < a$ (whereas $\pi_i(q_i, Q_{-i}) = -cq_i$ if $q_i + Q_{-i} \geq a$). It is again true that the monopoly quantity $q_m = (a-c)/2$ strictly dominates any higher quantity. That is, for any $x > 0$, $\pi_i(q_m, Q_{-i}) > \pi_i(q_m + x, Q_{-i})$ for all $Q_{-i} \geq 0$, just as in the first step in the duopoly case. Since

---

firm i's best response when firm i is uncertain about q_j. Suppose firm i is uncertain about q_j but believes that the expected value of q_j is $E(q_j)$. Because $\pi_i(q_i, q_j)$ is linear in q_j, firm i's best response when it is uncertain in this way simply equals its best response when it is certain that firm j will choose $E(q_j)$—a case covered in the text.

there are two firms other than firm i, however, all we can say about Q_{-i} is that it is between zero and $a - c$, because q_j and q_k are between zero and $(a - c)/2$. But this implies that no quantity $q_i \geq 0$ is strictly dominated for firm i, because for each q_i between zero and $(a - c)/2$ there exists a value of Q_{-i} between zero and $a-c$ (namely, $Q_{-i} = a-c-2q_i$) such that q_i is firm i's best response to Q_{-i}. Thus, no further strategies can be eliminated.

1.2.B Bertrand Model of Duopoly

We next consider a different model of how two duopolists might interact, based on Bertrand's (1883) suggestion that firms actually choose prices, rather than quantities as in Cournot's model. It is important to note that Bertrand's model is a *different game* than Cournot's model: the strategy spaces are different, the payoff functions are different, and (as will become clear) the behavior in the Nash equilibria of the two models is different. Some authors summarize these differences by referring to the Cournot and Bertrand equilibria. Such usage may be misleading: it refers to the difference between the Cournot and Bertrand games, and to the difference between the equilibrium behavior in these games, *not* to a difference in the equilibrium concept used in the games. *In both games, the equilibrium concept used is the Nash equilibrium defined in the previous section.*

We consider the case of differentiated products. (See Problem 1.7 for the case of homogeneous products.) If firms 1 and 2 choose prices p_1 and p_2, respectively, the quantity that consumers demand from firm i is

$$q_i(p_i, p_j) = a - p_i + bp_j,$$

where $b > 0$ reflects the extent to which firm i's product is a substitute for firm j's product. (This is an unrealistic demand function because demand for firm i's product is positive even when firm i charges an arbitrarily high price, provided firm j also charges a high enough price. As will become clear, the problem makes sense only if $b < 2$.) As in our discussion of the Cournot model, we assume that there are no fixed costs of production and that marginal costs are constant at c, where $c < a$, and that the firms act (i.e., choose their prices) simultaneously.

As before, the first task in the process of finding the Nash equilibrium is to translate the problem into a normal-form game. There are again two players. This time, however, the strategies available to each firm are the different prices it might charge, rather than the different quantities it might produce. We will assume that negative prices are not feasible but that any nonnegative price can be charged—there is no restriction to prices denominated in pennies, for instance. Thus, each firm's strategy space can again be represented as $S_i = [0, \infty)$, the nonnegative real numbers, and a typical strategy s_i is now a price choice, $p_i \geq 0$.

We will again assume that the payoff function for each firm is just its profit. The profit to firm i when it chooses the price p_i and its rival chooses the price p_j is

$$\pi_i(p_i, p_j) = q_i(p_i, p_j)[p_i - c] = [a - p_i + bp_j][p_i - c].$$

Thus, the price pair (p_1^*, p_2^*) is a Nash equilibrium if, for each firm i, p_i^* solves

$$\max_{0 \leq p_i < \infty} \pi_i(p_i, p_j^*) = \max_{0 \leq p_i < \infty} [a - p_i + bp_j^*][p_i - c].$$

The solution to firm i's optimization problem is

$$p_i^* = \frac{1}{2}(a + bp_j^* + c).$$

Therefore, if the price pair (p_1^*, p_2^*) is to be a Nash equilibrium, the firms' price choices must satisfy

$$p_1^* = \frac{1}{2}(a + bp_2^* + c)$$

and

$$p_2^* = \frac{1}{2}(a + bp_1^* + c).$$

Solving this pair of equations yields

$$p_1^* = p_2^* = \frac{a + c}{2 - b}.$$

1.2.C Final-Offer Arbitration

Many public-sector workers are forbidden to strike; instead, wage disputes are settled by binding arbitration. (Major league baseball may be a higher-profile example than the public sector but is substantially less important economically.) Many other disputes, including medical malpractice cases and claims by shareholders against their stockbrokers, also involve arbitration. The two major forms of arbitration are *conventional* and *final-offer* arbitration. In final-offer arbitration, the two sides make wage offers and then the arbitrator picks one of the offers as the settlement. In conventional arbitration, in contrast, the arbitrator is free to impose any wage as the settlement. We now derive the Nash equilibrium wage offers in a model of final-offer arbitration developed by Farber (1980).[9]

Suppose the parties to the dispute are a firm and a union and the dispute concerns wages. Let the timing of the game be as follows. First, the firm and the union simultaneously make offers, denoted by w_f and w_u, respectively. Second, the arbitrator chooses one of the two offers as the settlement. (As in many so-called static games, this is really a dynamic game of the kind to be discussed in Chapter 2, but here we reduce it to a static game between the firm and the union by making assumptions about the arbitrator's behavior in the second stage.) Assume that the arbitrator has an ideal settlement she would like to impose, denoted by x. Assume further that, after observing the parties' offers, w_f and w_u, the arbitrator simply chooses the offer that is closer to x: provided that $w_f < w_u$ (as is intuitive, and will be shown to be true), the arbitrator chooses w_f if $x < (w_f + w_u)/2$ and chooses w_u if $x > (w_f + w_u)/2$; see Figure 1.2.3. (It will be immaterial what happens if $x = (w_f + w_u)/2$. Suppose the arbitrator flips a coin.)

The arbitrator knows x but the parties do not. The parties believe that x is randomly distributed according to a cumulative probability distribution denoted by $F(x)$, with associated probability density function denoted by $f(x)$.[10] Given our specification of the arbitrator's behavior, if the offers are w_f and w_u

---

[9] This application involves some basic concepts in probability: a cumulative probability distribution, a probability density function, and an expected value. Terse definitions are given as needed; for more detail, consult any introductory probability text.

[10] That is, the probability that x is less than an arbitrary value x^* is denoted $F(x^*)$, and the derivative of this probability with respect to x^* is denoted $f(x^*)$. Since $F(x^*)$ is a probability, we have $0 \leq F(x^*) \leq 1$ for any x^*. Furthermore, if $x^{**} > x^*$ then $F(x^{**}) \geq F(x^*)$, so $f(x^*) \geq 0$ for every x^*.

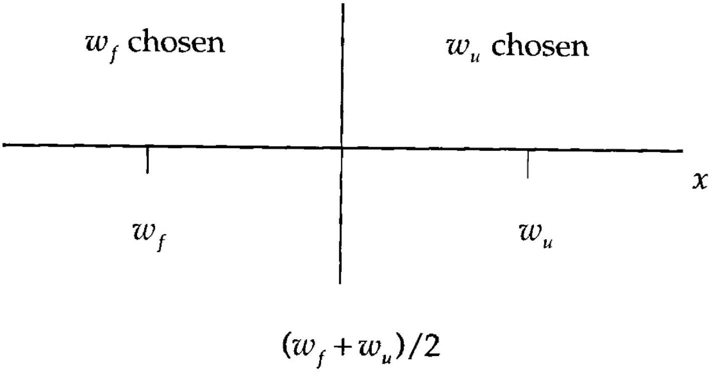

Figure 1.2.3.

then the parties believe that the probabilities Prob$\{w_f$ chosen$\}$ and Prob$\{w_u$ chosen$\}$ can be expressed as

$$\text{Prob}\{w_f \text{ chosen}\} = \text{Prob}\left\{x < \frac{w_f + w_u}{2}\right\} = F\left(\frac{w_f + w_u}{2}\right)$$

and

$$\text{Prob}\{w_u \text{ chosen}\} = 1 - F\left(\frac{w_f + w_u}{2}\right).$$

Thus, the expected wage settlement is

$$\begin{aligned} w_f &\cdot \text{Prob}\{w_f \text{ chosen}\} + w_u \cdot \text{Prob}\{w_u \text{ chosen}\} \\ &= w_f \cdot F\left(\frac{w_f + w_u}{2}\right) + w_u \cdot \left[1 - F\left(\frac{w_f + w_u}{2}\right)\right]. \end{aligned}$$

We assume that the firm wants to minimize the expected wage settlement imposed by the arbitrator and the union wants to maximize it.

If the pair of offers (w_f^*, w_u^*) is to be a Nash equilibrium of the game between the firm and the union, w_f^* must solve[11]

$$\min_{w_f} \; w_f \cdot F\left(\frac{w_f + w_u^*}{2}\right) + w_u^* \cdot \left[1 - F\left(\frac{w_f + w_u^*}{2}\right)\right]$$

and w_u^* must solve

$$\max_{w_u} \; w_f^* \cdot F\left(\frac{w_f^* + w_u}{2}\right) + w_u \cdot \left[1 - F\left(\frac{w_f^* + w_u}{2}\right)\right].$$

Thus, the wage-offer pair (w_f^*, w_u^*) must solve the first-order conditions for these optimization problems,

$$(w_u^* - w_f^*) \cdot \frac{1}{2} f\left(\frac{w_f^* + w_u^*}{2}\right) = F\left(\frac{w_f^* + w_u^*}{2}\right)$$

and

$$(w_u^* - w_f^*) \cdot \frac{1}{2} f\left(\frac{w_f^* + w_u^*}{2}\right) = \left[1 - F\left(\frac{w_f^* + w_u^*}{2}\right)\right].$$

(We defer considering whether these first-order conditions are sufficient.) Since the left-hand sides of these first-order conditions are equal, the right-hand sides must also be equal, which implies that

$$F\left(\frac{w_f^* + w_u^*}{2}\right) = \frac{1}{2}; \tag{1.2.2}$$

that is, the average of the offers must equal the median of the arbitrator's preferred settlement. Substituting (1.2.2) into either of the first-order conditions then yields

$$w_u^* - w_f^* = \frac{1}{f\left(\frac{w_f^* + w_u^*}{2}\right)}; \tag{1.2.3}$$

that is, the gap between the offers must equal the reciprocal of the value of the density function at the median of the arbitrator's preferred settlement.

---

[11] In formulating the firm's and the union's optimization problems, we have assumed that the firm's offer is less than the union's offer. It is straightforward to show that this inequality must hold in equilibrium.

In order to produce an intuitively appealing comparative-static result, we now consider an example. Suppose the arbitrator's preferred settlement is normally distributed with mean m and variance σ^2, in which case the density function is

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{1}{2\sigma^2}(x-m)^2\right\}.$$

(In this example, one can show that the first-order conditions given earlier are sufficient.) Because a normal distribution is symmetric around its mean, the median of the distribution equals the mean of the distribution, m. Therefore, (1.2.2) becomes

$$\frac{w_f^* + w_u^*}{2} = m$$

and (1.2.3) becomes

$$w_u^* - w_f^* = \frac{1}{f(m)} = \sqrt{2\pi\sigma^2},$$

so the Nash equilibrium offers are

$$w_u^* = m + \sqrt{\frac{\pi\sigma^2}{2}} \quad \text{and} \quad w_f^* = m - \sqrt{\frac{\pi\sigma^2}{2}}.$$

Thus, in equilibrium, the parties' offers are centered around the expectation of the arbitrator's preferred settlement (i.e., m), and the gap between the offers increases with the parties' uncertainty about the arbitrator's preferred settlement (i.e., σ^2).

The intuition behind this equilibrium is simple. Each party faces a trade-off. A more aggressive offer (i.e., a lower offer by the firm or a higher offer by the union) yields a better payoff if it is chosen as the settlement by the arbitrator but is less likely to be chosen. (We will see in Chapter 3 that a similar trade-off arises in a first-price, sealed-bid auction: a lower bid yields a better payoff if it is the winning bid but reduces the chances of winning.) When there is more uncertainty about the arbitrator's preferred settlement (i.e., σ^2 is higher), the parties can afford to be more aggressive because an aggressive offer is less likely to be wildly at odds with the arbitrator's preferred settlement. When there is hardly any uncertainty, in contrast, neither party can afford to make an offer far from the mean because the arbitrator is very likely to prefer settlements close to m.

1.2.D The Problem of the Commons

Since at least Hume (1739), political philosophers and economists have understood that if citizens respond only to private incentives, public goods will be underprovided and public resources overutilized. Today, even a casual inspection of the earth's environment reveals the force of this idea. Hardin's (1968) much cited paper brought the problem to the attention of noneconomists. Here we analyze a bucolic example.

Consider the n farmers in a village. Each summer, all the farmers graze their goats on the village green. Denote the number of goats the i^{th} farmer owns by g_i and the total number of goats in the village by $G = g_1 + \cdots + g_n$. The cost of buying and caring for a goat is c, independent of how many goats a farmer owns. The value to a farmer of grazing a goat on the green when a total of G goats are grazing is $v(G)$ *per goat*. Since a goat needs at least a certain amount of grass in order to survive, there is a maximum number of goats that can be grazed on the green, $G_{\max}$: $v(G) > 0$ for $G < G_{\max}$ but $v(G) = 0$ for $G \geq G_{\max}$. Also, since the first few goats have plenty of room to graze, adding one more does little harm to those already grazing, but when so many goats are grazing that they are all just barely surviving (i.e., G is just below $G_{\max}$), then adding one more dramatically harms the rest. Formally: for $G < G_{\max}, v'(G) < 0$ and $v''(G) < 0$, as in Figure 1.2.4.

During the spring, the farmers simultaneously choose how many goats to own. Assume goats are continuously divisible. A strategy for farmer i is the choice of a number of goats to graze on the village green, g_i. Assuming that the strategy space is $[0, \infty)$ covers all the choices that could possibly be of interest to the farmer; $[0, G_{\max})$ would also suffice. The payoff to farmer i from grazing g_i goats when the numbers of goats grazed by the other farmers are $(g_1, \ldots, g_{i-1}, g_{i+1}, \ldots, g_n)$ is

$$g_i v(g_1 + \cdots + g_{i-1} + g_i + g_{i+1} + \cdots + g_n) - c\, g_i. \qquad (1.2.4)$$

Thus, if $(g_1^*, \ldots, g_n^*)$ is to be a Nash equilibrium then, for each i, g_i^* must maximize (1.2.4) given that the other farmers choose $(g_1^*, \ldots, g_{i-1}^*, g_{i+1}^*, \ldots, g_n^*)$. The first-order condition for this optimization problem is

$$v(g_i + g_{-i}^*) + g_i v'(g_i + g_{-i}^*) - c = 0, \qquad (1.2.5)$$

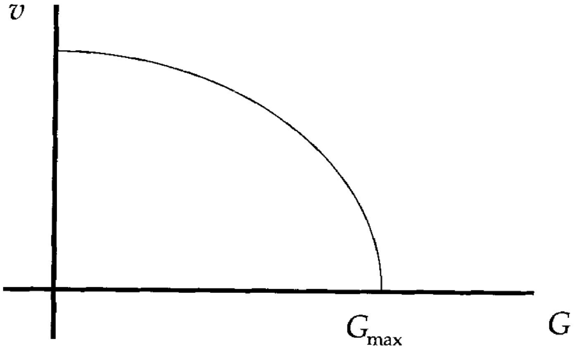

Figure 1.2.4.

where g^*_{-i} denotes $g_1^* + \cdots + g_{i-1}^* + g_{i+1}^* + \cdots + g_n^*$. Substituting g_i^* into (1.2.5), summing over all n farmers' first-order conditions, and then dividing by n yields

$$v(G^*) + \frac{1}{n}G^*v'(G^*) - c = 0, \tag{1.2.6}$$

where G^* denotes $g_1^* + \cdots + g_n^*$. In contrast, the social optimum, denoted by G^{**}, solves

$$\max_{0 \le G < \infty} Gv(G) - Gc,$$

the first-order condition for which is

$$v(G^{**}) + G^{**}v'(G^{**}) - c = 0. \tag{1.2.7}$$

Comparing (1.2.6) to (1.2.7) shows[12] that $G^* > G^{**}$: too many goats are grazed in the Nash equilibrium, compared to the social optimum. The first-order condition (1.2.5) reflects the incentives faced by a farmer who is already grazing g_i goats but is considering adding one more (or, strictly speaking, a tiny fraction of one more). The value of the additional goat is $v(g_i + g^*_{-i})$ and its cost is c. The harm to the farmer's existing goats is $v'(g_i+g^*_{-i})$ per goat, or $g_i v'(g_i + g^*_{-i})$ in total. The common resource is overutilized because each farmer considers only his or her own incentives, not the effect of his or her actions on the other farmers, hence the presence of $G^*v'(G^*)/n$ in (1.2.6) but $G^{**}v'(G^{**})$ in (1.2.7).

---

[12] Suppose, to the contrary, that $G^* \le G^{**}$. Then $v(G^*) \ge v(G^{**})$, since $v' < 0$. Likewise, $0 > v'(G^*) \ge v'(G^{**})$, since $v'' < 0$. Finally, $G^*/n < G^{**}$. Thus, the left-hand side of (1.2.6) strictly exceeds the left-hand side of (1.2.7), which is impossible since both equal zero.

1.3 Advanced Theory: Mixed Strategies and Existence of Equilibrium

1.3.A Mixed Strategies

In Section 1.1.C we defined S_i to be the set of strategies available to player i, and the combination of strategies $(s_1^*, \ldots, s_n^*)$ to be a Nash equilibrium if, for each player i, s_i^* is player i's best response to the strategies of the $n-1$ other players:

$$u_i(s_1^*, \ldots, s_{i-1}^*, s_i^*, s_{i+1}^*, \ldots, s_n^*) \geq u_i(s_1^*, \ldots, s_{i-1}^*, s_i, s_{i+1}^*, \ldots, s_n^*) \tag{NE}$$

for every strategy s_i in S_i. By this definition, there is no Nash equilibrium in the following game, known as *Matching Pennies*.

		Player 2	
		Heads	Tails
Player 1	Heads	−1, 1	1, −1
	Tails	1, −1	−1, 1

Matching Pennies

In this game, each player's strategy space is {Heads, Tails}. As a story to accompany the payoffs in the bi-matrix, imagine that each player has a penny and must choose whether to display it with heads or tails facing up. If the two pennies match (i.e., both are heads up or both are tails up) then player 2 wins player 1's penny; if the pennies do not match then 1 wins 2's penny. No pair of strategies can satisfy (NE), since if the players' strategies match—(Heads, Heads) or (Tails, Tails)—then player 1 prefers to switch strategies, while if the strategies do not match—(Heads, Tails) or (Tails, Heads)—then player 2 prefers to do so.

The distinguishing feature of Matching Pennies is that each player would like to outguess the other. Versions of this game also arise in poker, baseball, battle, and other settings. In poker, the analogous question is how often to bluff: if player i is known never to bluff then i's opponents will fold whenever i bids aggressively, thereby making it worthwhile for i to bluff on occasion; on the other hand, bluffing too often is also a losing strategy. In baseball, suppose that a pitcher can throw either a fastball or a curve and that a batter can hit either pitch if (but only if) it is anticipated correctly. Similarly, in battle, suppose that the attackers can choose between two locations (or two routes, such as "by land or by sea") and that the defense can parry either attack if (but only if) it is anticipated correctly.

In any game in which each player would like to outguess the other(s), there is no Nash equilibrium (at least as this equilibrium concept was defined in Section 1.1.C) because the solution to such a game necessarily involves uncertainty about what the players will do. We now introduce the notion of a *mixed strategy*, which we will interpret in terms of one player's uncertainty about what another player will do. (This interpretation was advanced by Harsanyi [1973]; we discuss it further in Section 3.2.A.) In the next section we will extend the definition of Nash equilibrium to include mixed strategies, thereby capturing the uncertainty inherent in the solution to games such as Matching Pennies, poker, baseball, and battle.

Formally, a mixed strategy for player i is a probability distribution over (some or all of) the strategies in S_i. We will hereafter refer to the strategies in S_i as player i's *pure strategies*. In the simultaneous-move games of complete information analyzed in this chapter, a player's pure strategies are the different actions the player could take. In Matching Pennies, for example, S_i consists of the two pure strategies Heads and Tails, so a mixed strategy for player i is the probability distribution $(q, 1-q)$, where q is the probability of playing Heads, $1-q$ is the probability of playing Tails, and $0 \leq q \leq 1$. The mixed strategy $(0,1)$ is simply the pure strategy Tails; likewise, the mixed strategy $(1,0)$ is the pure strategy Heads.

As a second example of a mixed strategy, recall Figure 1.1.1, where player 2 has the pure strategies Left, Middle, and Right. Here a mixed strategy for player 2 is the probability distribution $(q, r, 1-q-r)$, where q is the probability of playing Left, r is the probability of playing Middle, and $1-q-r$ is the probability of playing Right. As before, $0 \leq q \leq 1$, and now also $0 \leq r \leq 1$ and $0 \leq q+r \leq 1$. In this game, the mixed strategy $(1/3, 1/3, 1/3)$ puts equal probability on Left, Middle, and Right, whereas $(1/2, 1/2, 0)$ puts equal probability on Left and Middle but no probability on Right. As always, a player's pure strategies are simply the limiting cases of the player's mixed strategies—here player 2's pure strategy Left is the mixed strategy $(1, 0. 0)$, for example.

More generally, suppose that player i has K pure strategies: $S_i = \{s_{i1}, \ldots, s_{iK}\}$. Then a mixed strategy for player i is a probability distribution $(p_{i1}, \ldots. p_{iK})$, where p_{ik} is the probability that player i will play strategy s_{ik}, for $k = 1, \ldots, K$. Since p_{ik} is a probability, we require $0 \leq p_{ik} \leq 1$ for $k = 1, \ldots, K$ and $p_{i1} + \cdots + p_{iK} = 1$. We will use p_i to denote an arbitrary mixed strategy from the set of probability distributions over S_i, just as we use s_i to denote an arbitrary pure strategy from S_i.

Definition *In the normal-form game $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$, suppose $S_i = \{s_{i1}, \ldots, s_{iK}\}$. Then a **mixed strategy** for player i is a probability distribution $p_i = (p_{i1}, \ldots, p_{iK})$, where $0 \leq p_{ik} \leq 1$ for $k = 1, \ldots, K$ and $p_{i1} + \cdots + p_{iK} = 1$.*

We conclude this section by returning (briefly) to the notion of strictly dominated strategies introduced in Section 1.1.B, so as to illustrate the potential roles for mixed strategies in the arguments made there. Recall that if a strategy s_i is strictly dominated then there is no belief that player i could hold (about the strategies the other players will choose) such that it would be optimal to play s_i. The converse is also true, provided we allow for mixed strategies: if there is no belief that player i could hold (about the strategies the other players will choose) such that it would be optimal to play the strategy s_i, then there exists another strategy that strictly dominates s_i.[13] The games in Figures 1.3.1 and 1.3.2

---

[13] Pearce (1984) proves this result for the two-player case and notes that it holds for the n-player case provided that the players' mixed strategies are allowed to be correlated—that is, player i's belief about what player j will do must be allowed to be correlated with i's belief about what player k will do. Aumann (1987)

Player 2

		L	R
Player 1	T	3, —	0, —
	M	0, —	3, —
	B	1, —	1, —

Figure 1.3.1.

show that this converse would be false if we restricted attention to pure strategies.

Figure 1.3.1 shows that a given pure strategy may be strictly dominated by a mixed strategy, even if the pure strategy is not strictly dominated by any other pure strategy. In this game, for any belief $(q, 1-q)$ that player 1 could hold about 2's play, 1's best response is either T (if $q \geq 1/2$) or M (if $q \leq 1/2$), but never B. Yet B is not strictly dominated by either T or M. The key is that B is strictly dominated by a mixed strategy: if player 1 plays T with probability 1/2 and M with probability 1/2 then 1's expected payoff is 3/2 no matter what (pure or mixed) strategy 2 plays, and 3/2 exceeds the payoff of 1 that playing B surely produces. This example illustrates the role of mixed strategies in finding "another strategy that strictly dominates s_i."

Player 2

		L	R
Player 1	T	3, —	0, —
	M	0, —	3, —
	B	2, —	2, —

Figure 1.3.2.

suggests that such correlation in i's beliefs is entirely natural, even if j and k make their choices completely independently: for example, i may know that both j and k went to business school, or perhaps to the same business school, but may not know what is taught there.

Figure 1.3.2 shows that a given pure strategy can be a best response to a mixed strategy, even if the pure strategy is not a best response to any other pure strategy. In this game, B is not a best response for player 1 to either L or R by player 2, but B is the best response for player 1 to the mixed strategy $(q, 1-q)$ by player 2, provided $1/3 < q < 2/3$. This example illustrates the role of mixed strategies in the "belief that player i could hold."

1.3.B Existence of Nash Equilibrium

In this section we discuss several topics related to the existence of Nash equilibrium. First, we extend the definition of Nash equilibrium given in Section 1.1.C to allow for mixed strategies. Second, we apply this extended definition to Matching Pennies and the Battle of the Sexes. Third, we use a graphical argument to show that any two-player game in which each player has two pure strategies has a Nash equilibrium (possibly involving mixed strategies). Finally, we state and discuss Nash's (1950) Theorem, which guarantees that any finite game (i.e., any game with a finite number of players, each of whom has a finite number of pure strategies) has a Nash equilibrium (again, possibly involving mixed strategies).

Recall that the definition of Nash equilibrium given in Section 1.1.C guarantees that each player's pure strategy is a best response to the other players' pure strategies. To extend the definition to include mixed strategies, we simply require that each player's mixed strategy be a best response to the other players' mixed strategies. Since any pure strategy can be represented as the mixed strategy that puts zero probability on all of the player's other pure strategies, this extended definition subsumes the earlier one.

Computing player i's best response to a mixed strategy by player j illustrates the interpretation of player j's mixed strategy as representing player i's uncertainty about what player j will do. We begin with Matching Pennies as an example. Suppose that player 1 believes that player 2 will play Heads with probability q and Tails with probability $1-q$; that is, 1 believes that 2 will play the mixed strategy $(q, 1-q)$. Given this belief, player 1's expected payoffs are $q \cdot (-1) + (1-q) \cdot 1 = 1 - 2q$ from playing Heads and $q \cdot 1 + (1-q) \cdot (-1) = 2q - 1$ from playing Tails. Since $1 - 2q > 2q - 1$ if and only if $q < 1/2$, player 1's best pure-strategy response is

Heads if $q < 1/2$ and Tails if $q > 1/2$, and player 1 is indifferent between Heads and Tails if $q = 1/2$. It remains to consider possible mixed-strategy responses by player 1.

Let $(r, 1-r)$ denote the mixed strategy in which player 1 plays Heads with probability r. For each value of q between zero and one, we now compute the value(s) of r, denoted $r^*(q)$, such that $(r, 1-r)$ is a best response for player 1 to $(q, 1-q)$ by player 2. The results are summarized in Figure 1.3.3. Player 1's expected payoff from playing $(r, 1-r)$ when 2 plays $(q, 1-q)$ is

$$\begin{aligned} rq \cdot (-1) + r(1-q) \cdot 1 + (1-r)q \cdot 1 + (1-r)(1-q) \cdot (-1) \\ = (2q-1) + r(2-4q), \qquad (1.3.1) \end{aligned}$$

where rq is the probability of (Heads, Heads), $r(1-q)$ the probability of (Heads, Tails), and so on.[14] Since player 1's expected payoff is increasing in r if $2 - 4q > 0$ and decreasing in r if $2 - 4q < 0$, player 1's best response is $r = 1$ (i.e., Heads) if $q < 1/2$ and $r = 0$ (i.e., Tails) if $q > 1/2$, as indicated by the two horizontal segments of $r^*(q)$ in Figure 1.3.3. This statement is stronger than the closely related statement in the previous paragraph: there we considered only pure strategies and found that if $q < 1/2$ then Heads is the best pure strategy and that if $q > 1/2$ then Tails is the best pure strategy; here we consider all pure and mixed strategies but again find that if $q < 1/2$ then Heads is the best of all (pure or mixed) strategies and that if $q > 1/2$ then Tails is the best of all strategies.

The nature of player 1's best response to $(q, 1-q)$ changes when $q = 1/2$. As noted earlier, when $q = 1/2$ player 1 is indifferent between the pure strategies Heads and Tails. Furthermore, because player 1's expected payoff in (1.3.1) is independent of r when $q = 1/2$, player 1 is also indifferent among all mixed strategies $(r, 1-r)$. That is, when $q = 1/2$ the mixed strategy $(r, 1-r)$

---

[14] The events A and B are *independent* if Prob{A and B} = Prob{A}·Prob{B}. Thus, in writing rq for the probability that 1 plays Heads and 2 plays Heads, we are assuming that 1 and 2 make their choices independently, as befits the description we gave of simultaneous-move games. See Aumann (1974) for the definition of *correlated equilibrium*, which applies to games in which the players' choices can be correlated (because the players observe the outcome of a random event, such as a coin flip, before choosing their strategies).

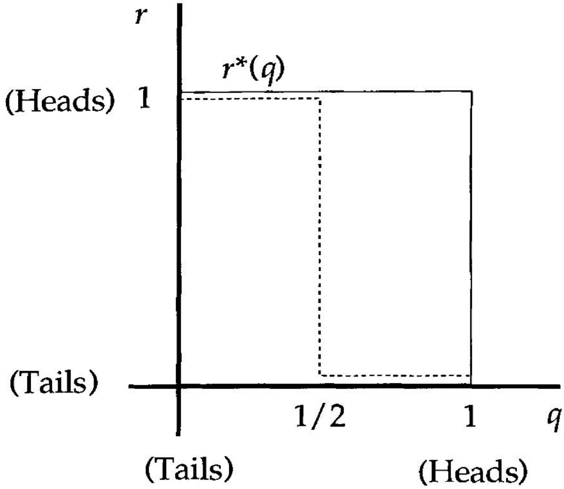

Figure 1.3.3.

is a best response to $(q, 1-q)$ for any value of r between zero and one. Thus, $r^*(1/2)$ is the entire interval $[0,1]$, as indicated by the vertical segment of $r^*(q)$ in Figure 1.3.3. In the analysis of the Cournot model in Section 1.2.A, we called $R_i(q_j)$ firm i's best-response *function*. Here, because there exists a value of q such that $r^*(q)$ has more than one value, we call $r^*(q)$ player 1's best-response *correspondence*.

To derive player i's best response to player j's mixed strategy more generally, and to give a formal statement of the extended definition of Nash equilibrium, we now restrict attention to the two-player case, which captures the main ideas as simply as possible. Let J denote the number of pure strategies in S_1 and K the number in S_2. We will write $S_1 = \{s_{11}, \ldots, s_{1J}\}$ and $S_2 = \{s_{21}, \ldots, s_{2K}\}$, and we will use s_{1j} and s_{2k} to denote arbitrary pure strategies from S_1 and S_2, respectively.

If player 1 believes that player 2 will play the strategies $(s_{21}, \ldots, s_{2K})$ with the probabilities $(p_{21}, \ldots, p_{2K})$ then player 1's expected payoff from playing the pure strategy s_{1j} is

$$\sum_{k=1}^{K} p_{2k} u_1(s_{1j}, s_{2k}), \tag{1.3.2}$$

and player 1's expected payoff from playing the mixed strategy $p_1 = (p_{11}, \ldots, p_{1J})$ is

$$\begin{aligned} v_1(p_1, p_2) &= \sum_{j=1}^{J} p_{1j} \left[\sum_{k=1}^{K} p_{2k} u_1(s_{1j}, s_{2k}) \right] \\ &= \sum_{j=1}^{J} \sum_{k=1}^{K} p_{1j} \cdot p_{2k} u_1(s_{1j}, s_{2k}), \end{aligned} \tag{1.3.3}$$

where $p_{1j} \cdot p_{2k}$ is the probability that 1 plays s_{1j} and 2 plays s_{2k}. Player 1's expected payoff from the mixed strategy p_1, given in (1.3.3), is the weighted sum of the expected payoff for each of the pure strategies $\{s_{11}, \ldots, s_{1J}\}$, given in (1.3.2), where the weights are the probabilities $(p_{11}, \ldots, p_{1J})$. Thus, for the mixed strategy $(p_{11}, \ldots, p_{1J})$ to be a best response for player 1 to 2's mixed strategy p_2, it must be that $p_{1j} > 0$ only if

$$\sum_{k=1}^{K} p_{2k} u_1(s_{1j}, s_{2k}) \geq \sum_{k=1}^{K} p_{2k} u_1(s_{1j'}, s_{2k})$$

for every $s_{1j'}$ in S_1. That is, for a mixed strategy to be a best response to p_2 it must put positive probability on a given pure strategy only if the pure strategy is itself a best response to p_2. Conversely, if player 1 has several pure strategies that are best responses to p_2, then any mixed strategy that puts all its probability on some or all of these pure-strategy best responses (and zero probability on all other pure strategies) is also a best response for player 1 to p_2.

To give a formal statement of the extended definition of Nash equilibrium, we need to compute player 2's expected payoff when players 1 and 2 play the mixed strategies p_1 and p_2 respectively. If player 2 believes that player 1 will play the strategies $(s_{11}, \ldots, s_{1J})$ with the probabilities $(p_{11}, \ldots, p_{1J})$, then player 2's expected payoff from playing the strategies $(s_{21}, \ldots, s_{2K})$ with the probabilities $(p_{21}, \ldots, p_{2K})$ is

$$\begin{aligned} v_2(p_1, p_2) &= \sum_{k=1}^{K} p_{2k} \left[\sum_{j=1}^{J} p_{1j} u_2(s_{1j}, s_{2k}) \right] \\ &= \sum_{j=1}^{J} \sum_{k=1}^{K} p_{1j} \cdot p_{2k} u_2(s_{1j}, s_{2k}). \end{aligned}$$

Given $v_1(p_1, p_2)$ and $v_2(p_1, p_2)$ we can restate the requirement of Nash equilibrium that each player's mixed strategy be a best response to the other player's mixed strategy: for the pair of mixed strategies (p_1^*, p_2^*) to be a Nash equilibrium, p_1^* must satisfy

$$v_1(p_1^*, p_2^*) \geq v_1(p_1, p_2^*) \tag{1.3.4}$$

for every probability distribution p_1 over S_1, and p_2^* must satisfy

$$v_2(p_1^*, p_2^*) \geq v_2(p_1^*, p_2) \tag{1.3.5}$$

for every probability distribution p_2 over S_2.

Definition *In the two-player normal-form game $G = \{S_1, S_2; u_1, u_2\}$, the mixed strategies (p_1^*, p_2^*) are a* ***Nash equilibrium*** *if each player's mixed strategy is a best response to the other player's mixed strategy: (1.3.4) and (1.3.5) must hold.*

We next apply this definition to Matching Pennies and the Battle of the Sexes. To do so, we use the graphical representation of player i's best response to player j's mixed strategy introduced in Figure 1.3.3. To complement Figure 1.3.3, we now compute the value(s) of q, denoted $q^*(r)$, such that $(q, 1-q)$ is a best response for player 2 to $(r, 1-r)$ by player 1. The results are summarized in Figure 1.3.4. If $r < 1/2$ then 2's best response is Tails, so $q^*(r) = 0$; likewise, if $r > 1/2$ then 2's best response is Heads, so $q^*(r) = 1$. If $r = 1/2$ then player 2 is indifferent not only between Heads and Tails but also among all the mixed strategies $(q, 1-q)$, so $q^*(1/2)$ is the entire interval $[0, 1]$.

After flipping and rotating Figure 1.3.4, we have Figure 1.3.5. Figure 1.3.5 is less convenient than Figure 1.3.4 as a representation

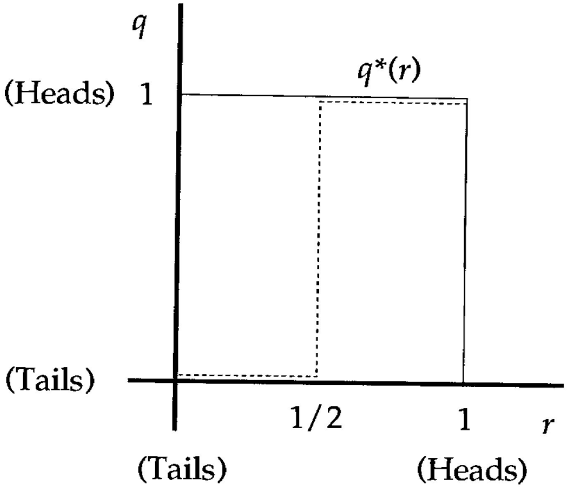

Figure 1.3.4.

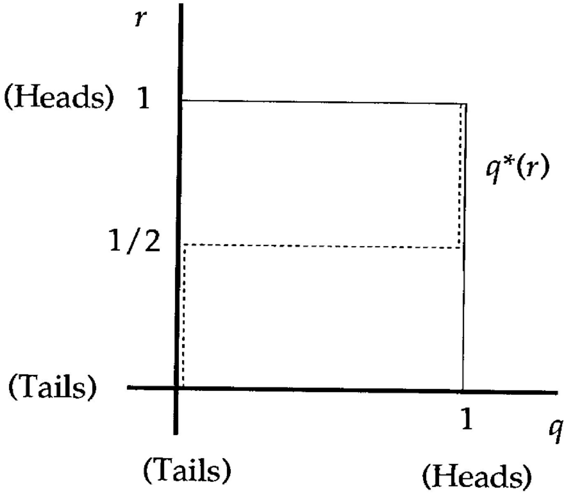

Figure 1.3.5.

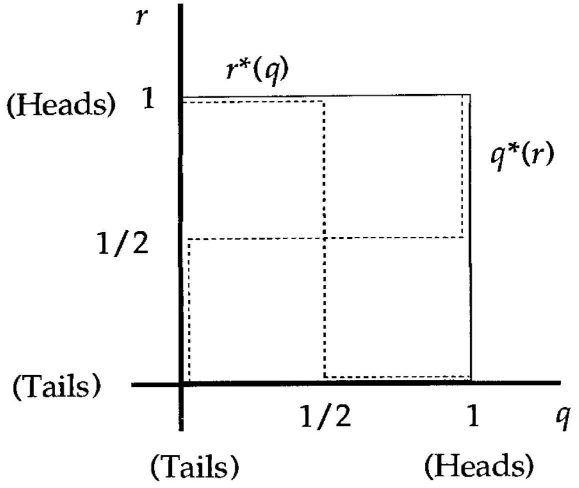

Figure 1.3.6.

of player 2's best response to player 1's mixed strategy, but it can be combined with Figure 1.3.3 to produce Figure 1.3.6.

Figure 1.3.6 is analogous to Figure 1.2.1 from the Cournot analysis in Section 1.2.A. Just as the intersection of the best-response functions $R_2(q_1)$ and $R_1(q_2)$ gave the Nash equilibrium of the Cournot game, the intersection of the best-response correspondences $r^*(q)$ and $q^*(r)$ yields the (mixed-strategy) Nash equilibrium in Matching Pennies: if player i plays $(1/2, 1/2)$ then $(1/2, 1/2)$ is a best response for player j, as required for Nash equilibrium.

It is worth emphasizing that such a mixed-strategy Nash equilibrium does *not* rely on any player flipping coins, rolling dice, or otherwise choosing a strategy at random. Rather, we interpret player j's mixed strategy as a statement of player i's uncertainty about player j's choice of a (pure) strategy. In baseball, for example, the pitcher might decide whether to throw a fastball or a curve based on how well each pitch was thrown during pregame practice. If the batter understands how the pitcher will make a choice but did not observe the pitcher's practice, then the batter may believe that the pitcher is equally likely to throw a fastball or a curve. We would then represent the batter's belief by the pitcher's mixed strategy $(1/2, 1/2)$, when in fact the pitcher chooses a pure strategy based on information unavailable to the batter. Stated more generally, the idea is to endow player j with a small amount of private information such that, depending on the realization of the private information, player j slightly prefers one of the relevant pure strategies. Since player i does not observe j's private information, however, i remains uncertain about j's choice, and we represent i's uncertainty by j's mixed strategy. We provide a more formal statement of this interpretation of a mixed strategy in Section 3.2.A.

As a second example of a mixed-strategy Nash equilibrium, consider the Battle of the Sexes from Section 1.1.C. Let $(q, 1-q)$ be the mixed strategy in which Pat plays Opera with probability q, and let $(r, 1-r)$ be the mixed strategy in which Chris plays Opera with probability r. If Pat plays $(q, 1-q)$ then Chris's expected payoffs are $q \cdot 2 + (1-q) \cdot 0 = 2q$ from playing Opera and $q \cdot 0 + (1-q) \cdot 1 = 1-q$ from playing Fight. Thus, if $q > 1/3$ then Chris's best response is Opera (i.e., $r = 1$), if $q < 1/3$ then Chris's best response is Fight (i.e., $r = 0$), and if $q = 1/3$ then any value of r is a best response. Similarly, if Chris plays $(r, 1-r)$ then Pat's expected payoffs are $r \cdot 1 + (1-r) \cdot 0 = r$ from playing Opera and $r \cdot 0 + (1-r) \cdot 2 = 2(1-r)$ from playing Fight. Thus, if $r > 2/3$ then Pat's best response is Opera (i.e., $q = 1$), if $r < 2/3$ then Pat's best response is Fight (i.e., $q = 0$), and if $r = 2/3$ then any value of q is a best response. As shown in Figure 1.3.7, the mixed strategies $(q, 1-q) = (1/3, 2/3)$ for Pat and $(r, 1-r) = (2/3, 1/3)$ for Chris are therefore a Nash equilibrium.

Unlike in Figure 1.3.6, where there was only one intersection of the players' best-response correspondences, there are three intersections of $r^*(q)$ and $q^*(r)$ in Figure 1.3.7: $(q = 0, r = 0)$ and $(q = 1, r = 1)$, as well as $(q = 1/3. r = 2/3)$. The other two intersections represent the pure-strategy Nash equilibria (Fight, Fight) and (Opera, Opera) described in Section 1.1.C.

In any game, a Nash equilibrium (involving pure or mixed strategies) appears as an intersection of the players' best-response correspondences, even when there are more than two players, and even when some or all of the players have more than two pure strategies. Unfortunately, the only games in which the players' best-response correspondences have simple graphical representations are two-player games in which each player has only two

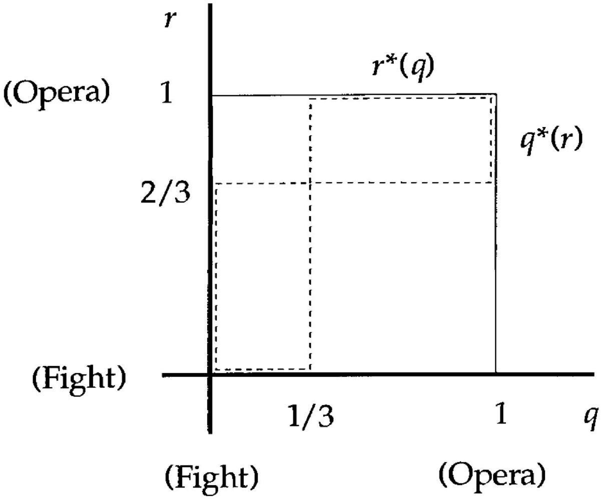

Figure 1.3.7.

		Player 2	
		Left	Right
Player 1	Up	x, —	y, —
	Down	z, —	w, —

Figure 1.3.8.

strategies. We turn next to a graphical argument that any such game has a Nash equilibrium (possibly involving mixed strategies).

Consider the payoffs for player 1 given in Figure 1.3.8. There are two important comparisons: x versus z, and y versus w. Based on these comparisons, we can define four main cases: (i) $x > z$ and $y > w$, (ii) $x < z$ and $y < w$, (iii) $x > z$ and $y < w$, and (iv) $x < z$ and $y > w$. We first discuss these four main cases, and then turn to the remaining cases involving $x = z$ or $y = w$.

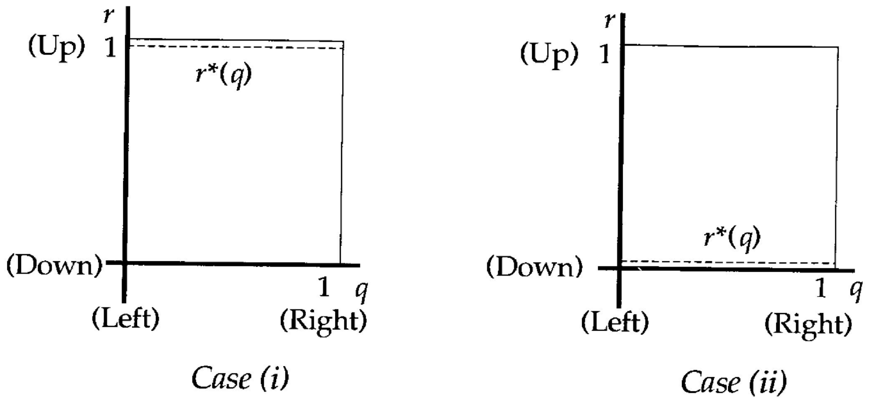

Figure 1.3.9.

In case (i) Up strictly dominates Down for player 1, and in case (ii) Down strictly dominates Up. Recall from the previous section that a strategy s_i is strictly dominated if and only if there is no belief that player i could hold (about the strategies the other players will choose) such that it would be optimal to play s_i. Thus, if $(q, 1-q)$ is a mixed strategy for player 2, where q is the probability that 2 will play Left, then in case (i) there is no value of q such that Down is optimal for player 1, and in case (ii) there is no value of q such that Up is optimal. Letting $(r, 1-r)$ denote a mixed strategy for player 1, where r is the probability that 1 will play Up, we can represent the best-response correspondences for cases (i) and (ii) as in Figure 1.3.9. (In these two cases the best-response correspondences are in fact best-response functions, since there is no value of q such that player 1 has multiple best responses.)

In cases (iii) and (iv), neither Up nor Down is strictly dominated. Thus, Up must be optimal for some values of q and Down optimal for others. Let $q' = (w-y)/(x-z+w-y)$. Then in case (iii) Up is optimal for $q > q'$ and Down for $q < q'$, whereas in case (iv) the reverse is true. In both cases, any value of r is optimal when $q = q'$. These best-response correspondences are given in Figure 1.3.10.

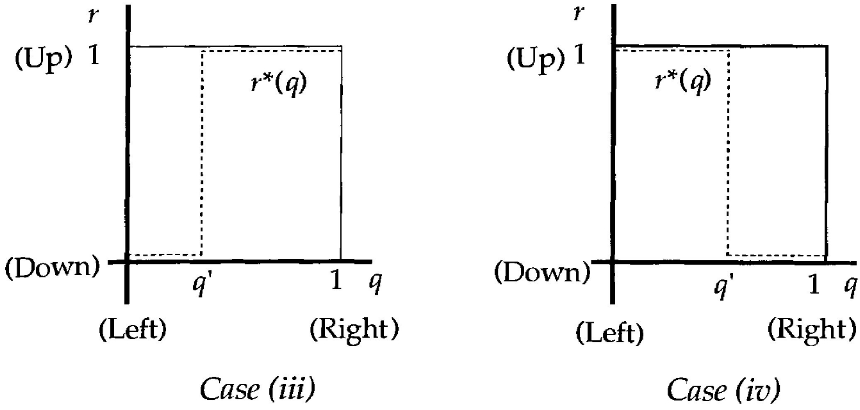

Figure 1.3.10.

Since $q' = 1$ if $x = z$ and $q' = 0$ if $y = w$, the best-response correspondences for cases involving either $x = z$ or $y = w$ are L-shaped (i.e., two adjacent sides of the unit square), as would occur in Figure 1.3.10 if $q' = 0$ or 1 in cases (iii) or (iv).

Adding arbitrary payoffs for player 2 to Figure 1.3.8 and performing the analogous computations yields the same four best-response correspondences, except that the horizontal axis measures r and the vertical q, as in Figure 1.3.4. Flipping and rotating these four figures, as was done to produce Figure 1.3.5, yields Figures 1.3.11 and 1.3.12. (In the latter figures, r' is defined analogously to q' in Figure 1.3.10.)

The crucial point is that given any of the four best-response correspondences for player 1, $r^*(q)$ from Figures 1.3.9 or 1.3.10, and any of the four for player 2, $q^*(r)$ from Figures 1.3.11 or 1.3.12, the pair of best-response correspondences has at least one intersection, so the game has at least one Nash equilibrium. Checking all sixteen possible pairs of best-response correspondences is left as an exercise. Instead, we describe the qualitative features that can result. There can be: (1) a single pure-strategy Nash equilibrium; (2) a single mixed-strategy equilibrium; or (3) two pure-strategy equilibria and a single mixed-strategy equilibrium. Recall from Figure 1.3.6 that Matching Pennies is an example of case (2), and from Figure 1.3.7 that the Battle of the Sexes is an example of

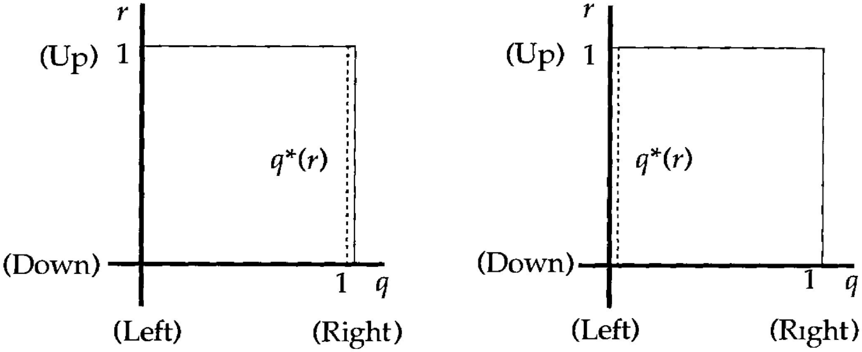
r
(Up) 1
q*(r)
(Down)
1 q
(Left)
(Right)
r
(Up) 1
q*(r)
(Down)
1 q
(Left)
(Right)

Case (i)

Case (ii)

Figure 1.3.11.

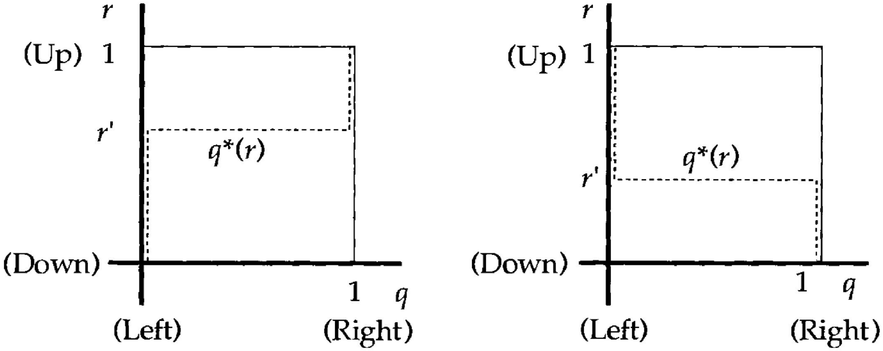
r
(Up) 1
r'
q*(r)
(Down)
1 q
(Left)
(Right)
r
(Up) 1
q*(r)
r'
(Down)
1 q
(Left)
(Right)

Case (iii)

Case (iv)

Figure 1.3.12.

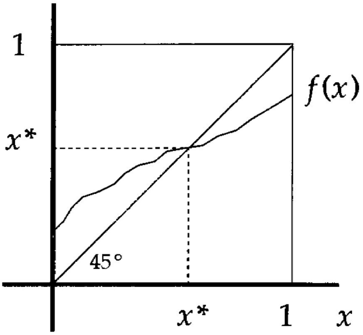

Figure 1.3.13.

case (3). The Prisoners' Dilemma is an example of case (1); it results from combining case (i) or (ii) of $r^*(q)$ with case (i) or (ii) or $q^*(r)$.[15]

We conclude this section with a discussion of the existence of a Nash equilibrium in more general games. If the above arguments for two-by-two games are stated mathematically rather than graphically, then they can be generalized to apply to n-player games with arbitrary finite strategy spaces.

Theorem (Nash 1950): *In the n-player normal-form game $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$, if n is finite and S_i is finite for every i then there exists at least one Nash equilibrium, possibly involving mixed strategies.*

The proof of Nash's Theorem involves a *fixed-point theorem*. As a simple example of a fixed-point theorem, suppose $f(x)$ is a continuous function with domain $[0,1]$ and range $[0,1]$. Then Brouwer's Fixed-Point Theorem guarantees that there exists at least one fixed point — that is, there exists at least one value x^* in $[0,1]$ such that $f(x^*) = x^*$. Figure 1.3.13 provides an example.

---

[15] The cases involving $x = z$ or $y = w$ do not violate the claim that the pair of best-response correspondences has at least one intersection. On the contrary, in addition to the qualitative features described in the text, there can now be two pure-strategy Nash equilibria without a mixed-strategy Nash equilibrium, and a continuum of mixed-strategy Nash equilibria.

Applying a fixed-point theorem to prove Nash's Theorem involves two steps: (1) showing that any fixed point of a certain correspondence is a Nash equilibrium; (2) using an appropriate fixed-point theorem to show that this correspondence must have a fixed point. The relevant correspondence is the n-player best-response correspondence. The relevant fixed-point theorem is due to Kakutani (1941), who generalized Brouwer's theorem to allow for (well-behaved) correspondences as well as functions.

The n-player best-response correspondence is computed from the n individual players' best-response correspondences as follows. Consider an arbitrary combination of mixed strategies $(p_1, \ldots, p_n)$. For each player i, derive i's best response(s) to the other players' mixed strategies $(p_1, \ldots, p_{i-1}, p_{i+1}, \ldots, p_n)$. Then construct the set of all possible combinations of one such best response for each player. (Formally, derive each player's best-response correspondence and then construct the cross-product of these n individual correspondences.) A combination of mixed strategies $(p_1^*, \ldots, p_n^*)$ is a fixed point of this correspondence if $(p_1^*, \ldots, p_n^*)$ belongs to the set of all possible combinations of the players' best responses to $(p_1^*, \ldots, p_n^*)$. That is, for each i, p_i^* must be (one of) player i's best response(s) to $(p_1^*, \ldots, p_{i-1}^*, p_{i+1}^*, \ldots, p_n^*)$, but this is precisely the statement that $(p_1^*, \ldots, p_n^*)$ is a Nash equilibrium. This completes step (1).

Step (2) involves the fact that each player's best-response correspondence is continuous, in an appropriate sense. The role of continuity in Brouwer's fixed-point theorem can be seen by modifying $f(x)$ in Figure 1.3.13: if $f(x)$ is discontinuous then it need not have a fixed point. In Figure 1.3.14, for example, $f(x) > x$ for all $x < x'$, but $f(x') < x'$ for $x \geq x'$.[16]

To illustrate the differences between $f(x)$ in Figure 1.3.14 and a player's best-response correspondence, consider Case (iii) in Figure 1.3.10: at $q = q'$, $r^*(q')$ includes zero, one, and the entire interval in between. (A bit more formally, $r^*(q')$ includes the limit of $r^*(q)$ as q approaches q' from the left, the limit of $r^*(q)$ as q approaches q' from the right, and all the values of r in between these two limits.) If $f(x')$ in Figure 1.3.14 behaved analogously to

---

[16] The value of $f(x')$ is indicated by the solid circle. The open circle indicates that $f(x')$ does not include this value. The dotted line is included only to indicate that both circles occur at $x = x'$; it does not indicate further values of $f(x')$.

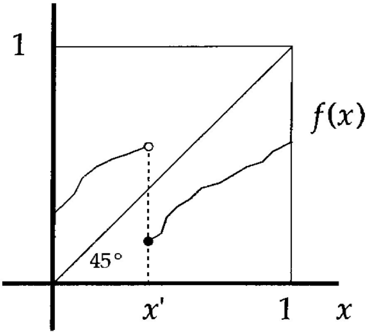

Figure 1.3.14.

player 1's best-response correspondence $r^*(q')$, then $f(x')$ would include not only the solid circle (as in the figure) but also the open circle and the entire interval in between, in which case $f(x)$ would have a fixed point at x'.

Each player's best-response correspondence always behaves the way $r^*(q')$ does in Figure 1.3.14: it always includes (the appropriate generalizations of) the limit from the left, the limit from the right, and all the values in between. The reason for this is that, as shown earlier for the two-player case, if player i has several pure strategies that are best responses to the other players' mixed strategies, then any mixed strategy p_i that puts all its probability on some or all of player i's pure-strategy best responses (and zero probability on all of player i's other pure strategies) is also a best response for player i. Because each player's best-response correspondence always behaves in this way, the n-player best-response correspondence does too; these properties satisfy the hypotheses of Kakutani's Theorem, so the latter correspondence has a fixed point.

Nash's Theorem guarantees that an equilibrium exists in a broad class of games, but none of the applications analyzed in Section 1.2 are members of this class (because each application has infinite strategy spaces). This shows that the hypotheses of Nash's Theorem are sufficient but not necessary conditions for an equilibrium to exist—there are many games that do not satisfy the hypotheses of the Theorem but nonetheless have one or more Nash equilibria.

1.4 Further Reading

On the assumptions underlying iterated elimination of strictly dominated strategies and Nash equilibrium, and on the interpretation of mixed strategies in terms of the players' beliefs, see Brandenburger (1992). On the relation between (Cournot-type) models where firms choose quantities and (Bertrand-type) models where firms choose prices, see Kreps and Scheinkman (1983), who show that in some circumstances the Cournot outcome occurs in a Bertrand-type model in which firms face capacity constraints (which they choose, at a cost, prior to choosing prices). On arbitration, see Gibbons (1988), who shows how the arbitrator's preferred settlement can depend on the information content of the parties' offers, in both final-offer and conventional arbitration. Finally, on the existence of Nash equilibrium, including pure-strategy equilibria in games with continuous strategy spaces, see Dasgupta and Maskin (1986).

1.5 Problems

Section 1.1

1.1. What is a game in normal form? What is a strictly dominated strategy in a normal-form game? What is a pure-strategy Nash equilibrium in a normal-form game?

1.2. In the following normal-form game, what strategies survive iterated elimination of strictly dominated strategies? What are the pure-strategy Nash equilibria?

	L	*C*	*R*
T	2.0	1,1	4,2
M	3,4	1,2	2,3
B	1,3	0,2	3,0

1.3. Players 1 and 2 are bargaining over how to split one dollar. Both players simultaneously name shares they would like to have, s_1 and s_2, where $0 \leq s_1, s_2 \leq 1$. If $s_1 + s_2 \leq 1$, then the players receive the shares they named; if $s_1 + s_2 > 1$, then both players receive zero. What are the pure-strategy Nash equilibria of this game?

Section 1.2

1.4. Suppose there are n firms in the Cournot oligopoly model. Let q_i denote the quantity produced by firm i, and let $Q = q_1 + \cdots + q_n$ denote the aggregate quantity on the market. Let P denote the market-clearing price and assume that inverse demand is given by $P(Q) = a - Q$ (assuming $Q < a$, else $P = 0$). Assume that the total cost of firm i from producing quantity q_i is $C_i(q_i) = cq_i$. That is, there are no fixed costs and the marginal cost is constant at c, where we assume $c < a$. Following Cournot, suppose that the firms choose their quantities simultaneously. What is the Nash equilibrium? What happens as n approaches infinity?

1.5. Consider the following two finite versions of the Cournot duopoly model. First, suppose each firm must choose either half the monopoly quantity, $q_m/2 = (a - c)/4$, or the Cournot equilibrium quantity, $q_c = (a - c)/3$. No other quantities are feasible. Show that this two-action game is equivalent to the Prisoners' Dilemma: each firm has a strictly dominated strategy, and both are worse off in equilibrium than they would be if they cooperated. Second, suppose each firm can choose either $q_m/2$, or q_c, or a third quantity, q'. Find a value for q' such that the game is equivalent to the Cournot model in Section 1.2.A, in the sense that (q_c, q_c) is a unique Nash equilibrium and both firms are worse off in equilibrium than they could be if they cooperated, but neither firm has a strictly dominated strategy.

1.6. Consider the Cournot duopoly model where inverse demand is $P(Q) = a - Q$ but firms have asymmetric marginal costs: c_1 for firm 1 and c_2 for firm 2. What is the Nash equilibrium if $0 < c_i < a/2$ for each firm? What if $c_1 < c_2 < a$ but $2c_2 > a + c_1$?

1.7. In Section 1.2.B, we analyzed the Bertrand duopoly model with differentiated products. The case of homogeneous products yields a stark conclusion. Suppose that the quantity that consumers demand from firm i is $a - p_i$ when $p_i < p_j$, 0 when $p_i > p_j$, and $(a - p_i)/2$ when $p_i = p_j$. Suppose also that there are no fixed costs and that marginal costs are constant at c, where $c < a$. Show that if the firms choose prices simultaneously, then the unique Nash equilibrium is that both firms charge the price c.

1.8. Consider a population of voters uniformly distributed along the ideological spectrum from left ($x = 0$) to right ($x = 1$). Each of the candidates for a single office simultaneously chooses a campaign platform (i.e., a point on the line between $x = 0$ and $x = 1$). The voters observe the candidates' choices, and then each voter votes for the candidate whose platform is closest to the voter's position on the spectrum. If there are two candidates and they choose platforms $x_1 = .3$ and $x_2 = .6$, for example, then all voters to the left of $x = .45$ vote for candidate 1, all those to the right vote for candidate 2, and candidate 2 wins the election with 55 percent of the vote. Suppose that the candidates care only about being elected—they do not really care about their platforms at all! If there are two candidates, what is the pure-strategy Nash equilibrium? If there are three candidates, exhibit a pure-strategy Nash equilibrium. (Assume that any candidates who choose the same platform equally split the votes cast for that platform, and that ties among the leading vote-getters are resolved by coin flips.) See Hotelling (1929) for an early model along these lines.

Section 1.3

1.9. What is a mixed strategy in a normal-form game? What is a mixed-strategy Nash equilibrium in a normal-form game?

1.10. Show that there are no mixed-strategy Nash equilibria in the three normal-form games analyzed in Section 1.1—the Prisoners' Dilemma, Figure 1.1.1, and Figure 1.1.4.

1.11. Solve for the mixed-strategy Nash equilibria in the game in Problem 1.2.

1.12. Find the mixed-strategy Nash equilibrium of the following normal-form game.

	L	R
T	2, 1	0, 2
B	1, 2	3, 0

1.13. Each of two firms has one job opening. Suppose that (for reasons not discussed here but relating to the value of filling each opening) the firms offer different wages: firm *i* offers the wage w_i, where $(1/2)w_1 < w_2 < 2w_1$. Imagine that there are two workers, each of whom can apply to only one firm. The workers simultaneously decide whether to apply to firm 1 or to firm 2. If only one worker applies to a given firm, that worker gets the job; if both workers apply to one firm, the firm hires one worker at random and the other worker is unemployed (which has a payoff of zero). Solve for the Nash equilibria of the workers' normal-form game. (For more on the wages the firms will choose, see Montgomery [1991].)

		Worker 2	
		Apply to Firm 1	Apply to Firm 2
Worker 1	Apply to Firm 1	$\frac{1}{2}w_1, \frac{1}{2}w_1$	w_1, w_2
	Apply to Firm 2	w_2, w_1	$\frac{1}{2}w_2, \frac{1}{2}w_2$

1.14. Show that Proposition B in Appendix 1.1.C holds for mixed- as well as pure-strategy Nash equilibria: the strategies played with positive probability in a mixed-strategy Nash equilibrium survive the process of iterated elimination of strictly dominated strategies.

1.6 References

Aumann, R. 1974. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics* 1:67–96.

———. 1976. "Agreeing to Disagree." *Annals of Statistics* 4:1236–39.

———. 1987. "Correlated Equilibrium as an Expression of Bayesian Rationality." *Econometrica* 55:1–18.

Bertrand, J. 1883. "Théorie Mathématique de la Richesse Sociale." *Journal des Savants* 499–508.

Brandenburger, A. 1992. "Knowledge and Equilibrium in Games." Forthcoming in *Journal of Economic Perspectives.*

Cournot, A. 1838. *Recherches sur les Principes Mathématiques de la Théorie des Richesses.* English edition: *Researches into the Mathematical Principles of the Theory of Wealth.* Edited by N. Bacon. New York: Macmillan, 1897.

Dasgupta, P., and E. Maskin. 1986. "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory." *Review of Economic Studies* 53:1–26.

Farber, H. 1980. "An Analysis of Final-Offer Arbitration." *Journal of Conflict Resolution* 35:683–705.

Friedman, J. 1971. "A Noncooperative Equilibrium for Supergames." *Review of Economic Studies* 28:1–12.

Gibbons, R. 1988. "Learning in Equilibrium Models of Arbitration." *American Economic Review* 78:896–912.

Hardin, G. 1968. "The Tragedy of the Commons." *Science* 162:1243–48.

Harsanyi, J. 1973. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points." *International Journal of Game Theory* 2:1–23.

Hotelling, H. 1929. "Stability in Competition." *Economic Journal* 39:41–57.

Hume, D. 1739. *A Treatise of Human Nature.* Reprint. London: J. M. Dent. 1952.

Kakutani, S. 1941. "A Generalization of Brouwer's Fixed Point Theorem." *Duke Mathematical Journal* 8:457–59.

Kreps, D., and J. Scheinkman. 1983. "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes." *Bell Journal of Economics* 14:326–37.

Montgomery, J. 1991. "Equilibrium Wage Dispersion and Interindustry Wage Differentials." *Quarterly Journal of Economics* 106: 163–79.

Nash, J. 1950. "Equilibrium Points in n-Person Games." *Proceedings of the National Academy of Sciences* 36:48–49.

Pearce, D. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52:1029–50.

Stackelberg, H. von. 1934. *Marktform und Gleichgewicht*. Vienna: Julius Springer.

Chapter 2

Dynamic Games of Complete Information

In this chapter we introduce dynamic games. We again restrict attention to games with complete information (i.e., games in which the players' payoff functions are common knowledge); see Chapter 3 for the introduction to games of incomplete information. In Section 2.1 we analyze dynamic games that have not only complete but also *perfect information,* by which we mean that at each move in the game the player with the move knows the full history of the play of the game thus far. In Sections 2.2 through 2.4 we consider games of complete but imperfect information: at some move the player with the move does not know the history of the game.

The central issue in all dynamic games is credibility. As an example of a noncredible threat, consider the following two-move game. First, player 1 chooses between giving player 2 \$1,000 and giving player 2 nothing. Second, player 2 observes player 1's move and then chooses whether or not to explode a grenade that will kill both players. Suppose player 2 threatens to explode the grenade unless player 1 pays the \$1,000. If player 1 believes the threat, then player 1's best response is to pay the \$1,000. But player 1 should not believe the threat, because it is noncredible: if player 2 were given the opportunity to carry out the threat, player 2 would choose not to carry it out. Thus, player 1 should pay player 2 nothing.[1]

In Section 2.1 we analyze the following class of dynamic games of complete and perfect information: first player 1 moves, then player 2 observes player 1's move, then player 2 moves and the game ends. The grenade game belongs to this class, as do Stackelberg's (1934) model of duopoly and Leontief's (1946) model of wage and employment determination in a unionized firm. We define the *backwards-induction outcome* of such games and briefly discuss its relation to Nash equilibrium (deferring the main discussion of this relation until Section 2.4). We solve for this outcome in the Stackelberg and Leontief models. We also derive the analogous outcome in Rubinstein's (1982) bargaining model, although this game has a potentially infinite sequence of moves and so does not belong to the above class of games.

In Section 2.2 we enrich the class of games analyzed in the previous section: first players 1 and 2 move simultaneously, then players 3 and 4 observe the moves chosen by 1 and 2, then players 3 and 4 move simultaneously and the game ends. As will be explained in Section 2.4, the simultaneity of moves here means that these games have imperfect information. We define the *subgame-perfect outcome* of such games, which is the natural extension of backwards induction to these games. We solve for this outcome in Diamond and Dybvig's (1983) model of bank runs, in a model of tariffs and imperfect international competition, and in Lazear and Rosen's (1981) model of tournaments.

In Section 2.3 we study *repeated games*, in which a fixed group of players plays a given game repeatedly, with the outcomes of all previous plays observed before the next play begins. The theme of the analysis is that (credible) threats and promises about future behavior can influence current behavior. We define *subgame-perfect Nash equilibrium* for repeated games and relate it to the backwards-induction and subgame-perfect outcomes defined in Sections 2.1 and 2.2. We state and prove the Folk Theorem for infinitely repeated games, and we analyze Friedman's (1971) model of collusion between Cournot duopolists, Shapiro and Stiglitz's (1984) model of efficiency wages, and Barro and Gordon's (1983) model of monetary policy.

---

[1] Player 1 might wonder whether an opponent who threatens to explode a grenade is crazy. We model such doubts as incomplete information—player 1 is unsure about player 2's payoff function. See Chapter 3.

In Section 2.4 we introduce the tools necessary to analyze a general dynamic game of complete information, whether with perfect or imperfect information. We define the *extensive-form* representation of a game and relate it to the normal-form representation introduced in Chapter 1. We also define subgame-perfect Nash equilibrium for general games. The main point (of both this section and the chapter as a whole) is that a dynamic game of complete information may have many Nash equilibria, but some of these may involve noncredible threats or promises. The subgame-perfect Nash equilibria are those that pass a credibility test.

2.1 Dynamic Games of Complete and Perfect Information

2.1.A Theory: Backwards Induction

The grenade game is a member of the following class of simple games of complete and perfect information:

1. Player 1 chooses an action a_1 from the feasible set A_1.
2. Player 2 observes a_1 and then chooses an action a_2 from the feasible set A_2.
3. Payoffs are $u_1(a_1, a_2)$ and $u_2(a_1, a_2)$.

Many economic problems fit this description.[2] Two examples

---

[2] Player 2's feasible set of actions, A_2, could be allowed to depend on player 1's action, a_1. Such dependence could be denoted by $A_2(a_1)$ or could be incorporated into player 2's payoff function, by setting $u_2(a_1, a_2) = -\infty$ for values of a_2 that are not feasible for a given a_1. Some moves by player 1 could even end the game, without player 2 getting a move; for such values of a_1, the set of feasible actions $A_2(a_1)$ contains only one element, so player 2 has no choice to make.

(discussed later in detail) are Stackelberg's model of duopoly and Leontief's model of wages and employment in a unionized firm. Other economic problems can be modeled by allowing for a longer sequence of actions, either by adding more players or by allowing players to move more than once. (Rubinstein's bargaining game, discussed in Section 2.1.D, is an example of the latter.) The key features of a dynamic game of complete and perfect information are that (i) the moves occur in sequence, (ii) all previous moves are observed before the next move is chosen, and (iii) the players' payoffs from each feasible combination of moves are common knowledge.

We solve a game from this class by backwards induction, as follows. When player 2 gets the move at the second stage of the game, he or she will face the following problem, given the action a_1 previously chosen by player 1:

$$\max_{a_2 \in A_2} u_2(a_1, a_2).$$

Assume that for each a_1 in A_1, player 2's optimization problem has a unique solution, denoted by $R_2(a_1)$. This is player 2's *reaction* (or best response) to player 1's action. Since player 1 can solve 2's problem as well as 2 can, player 1 should anticipate player 2's reaction to each action a_1 that 1 might take, so 1's problem at the first stage amounts to

$$\max_{a_1 \in A_1} u_1(a_1, R_2(a_1)).$$

Assume that this optimization problem for player 1 also has a unique solution, denoted by a_1^*. We will call $(a_1^*, R_2(a_1^*))$ the *backwards-induction outcome* of this game. The backwards-induction outcome does not involve noncredible threats: player 1 anticipates that player 2 will respond optimally to *any* action a_1 that 1 might choose, by playing $R_2(a_1)$; player 1 gives no credence to threats by player 2 to respond in ways that will not be in 2's self-interest when the second stage arrives.

Recall that in Chapter 1 we used the normal-form representation to study static games of complete information, and we focused on the notion of Nash equilibrium as a solution concept for such games. In this section's discussion of dynamic games, however, we have made no mention of either the normal-form representation or Nash equilibrium. Instead, we have given a verbal description of a game in (1)–(3), and we have defined the backwards-induction outcome as the solution to that game. In Section 2.4.A we will see that the verbal description in (1)–(3) is the extensive-form representation of the game. We will relate the extensive- and normal-form representations, but we will find that for dynamic games the extensive-form representation is often more convenient. In Section 2.4.B we will define subgame-perfect Nash equilibrium: a Nash equilibrium is subgame-perfect if it does not involve a noncredible threat, in a sense to be made precise. We will find that there may be multiple Nash equilibria in a game from the class defined by (1)–(3), but that the only subgame-perfect Nash equilibrium is the equilibrium associated with the backwards-induction outcome. This is an example of the observation in Section 1.1.C that some games have multiple Nash equilibria but have one equilibrium that stands out as the compelling solution to the game.

We conclude this section by exploring the rationality assumptions inherent in backwards-induction arguments. Consider the following three-move game, in which player 1 moves twice:

1. Player 1 chooses L or R, where L ends the game with payoffs of 2 to player 1 and 0 to player 2.

2. Player 2 observes 1's choice. If 1 chose R then 2 chooses L' or R', where L' ends the game with payoffs of 1 to both players.

3. Player 1 observes 2's choice (and recalls his or her own choice in the first stage). If the earlier choices were R and R' then 1 chooses L'' or R'', both of which end the game, L'' with payoffs of 3 to player 1 and 0 to player 2 and R'' with analogous payoffs of 0 and 2.

All these words can be translated into the following succinct game tree. (This is the extensive-form representation of the game, to be defined more generally in Section 2.4.) The top payoff in the pair of payoffs at the end of each branch of the game tree is player 1's, the bottom player 2's.

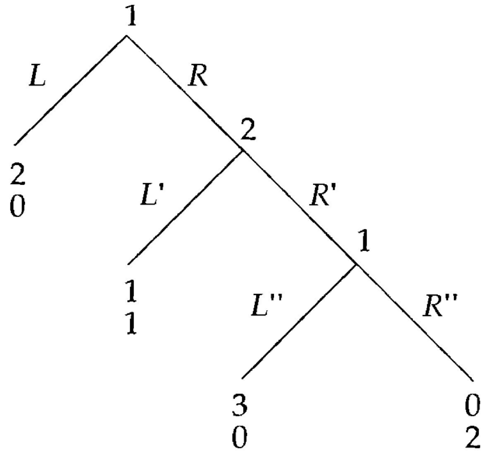

To compute the backwards-induction outcome of this game, we begin at the third stage (i.e., player 1's second move). Here player 1 faces a choice between a payoff of 3 from L'' and a payoff of 0 from R'', so L'' is optimal. Thus, at the second stage, player 2 anticipates that if the game reaches the third stage then 1 will play L'', which would yield a payoff of 0 for player 2. The second-stage choice for player 2 therefore is between a payoff of 1 from L' and a payoff of 0 from R', so L' is optimal. Thus, at the first stage, player 1 anticipates that if the game reaches the second stage then 2 will play L', which would yield a payoff of 1 for player 1. The first-stage choice for player 1 therefore is between a payoff of 2 from L and a payoff of 1 from R, so L is optimal.

This argument establishes that the backwards-induction outcome of this game is for player 1 to choose L in the first stage, thereby ending the game. Even though backwards induction predicts that the game will end in the first stage, an important part of the argument concerns what would happen if the game did not end in the first stage. In the second stage, for example, when player 2 anticipates that if the game reaches the third stage then 1 will play L'', 2 is assuming that 1 is rational. This assumption may seem inconsistent with the fact that 2 gets to move in the second stage only if 1 deviates from the backwards-induction outcome of the game. That is, it may seem that if 1 plays R in the first stage then 2 cannot assume in the second stage that 1 is rational, but this is not the case: if 1 plays R in the first stage then it cannot be common knowledge that both players are rational, but there remain reasons for 1 to have chosen R that do not contradict 2's assumption that 1 is rational.[3] One possibility is that it is common knowledge that player 1 is rational but not that player 2 is rational: if 1 thinks that 2 might not be rational, then 1 might choose R in the first stage, hoping that 2 will play R' in the second stage, thereby giving 1 the chance to play L'' in the third stage. Another possibility is that it is common knowledge that player 2 is rational but not that player 1 is rational: if 1 is rational but thinks that 2 thinks that 1 might not be rational, then 1 might choose R in the first stage, hoping that 2 will think that 1 is not rational and so play R' in the hope that 1 will play R'' in the third stage. Backwards induction assumes that 1's choice of R could be explained along these lines. For some games, however, it may be more reasonable to assume that 1 played R because 1 is indeed irrational. In such games, backwards induction loses much of its appeal as a prediction of play, just as Nash equilibrium does in games where game theory does not provide a unique solution and no convention will develop.

2.1.B Stackelberg Model of Duopoly

Stackelberg (1934) proposed a dynamic model of duopoly in which a dominant (or leader) firm moves first and a subordinate (or follower) firm moves second. At some points in the history of the U.S. automobile industry, for example, General Motors has seemed to play such a leadership role. (It is straightforward to extend what follows to allow for more than one following firm, such as Ford, Chrysler, and so on.) Following Stackelberg, we will develop the model under the assumption that the firms choose quantities, as in the Cournot model (where the firms' choices are simultaneous, rather than sequential as here). We leave it as an exercise to develop the analogous sequential-move model in which firms choose prices, as they do (simultaneously) in the Bertrand model.

The timing of the game is as follows: (1) firm 1 chooses a quantity $q_1 \geq 0$; (2) firm 2 observes q_1 and then chooses a quantity

---

[3] Recall from the discussion of iterated elimination of strictly dominated strategies (in Section 1.1.B) that it is common knowledge that the players are rational if all the players are rational, and all the players know that all the players are rational, and all the players know that all the players know that all the players are rational, and so on, ad infinitum.

$q_2 \geq 0$; (3) the payoff to firm i is given by the profit function

$$\pi_i(q_i, q_j) = q_i[P(Q) - c],$$

where $P(Q) = a - Q$ is the market-clearing price when the aggregate quantity on the market is $Q = q_1 + q_2$, and c is the constant marginal cost of production (fixed costs being zero).

To solve for the backwards-induction outcome of this game, we first compute firm 2's reaction to an arbitrary quantity by firm 1. $R_2(q_1)$ solves

$$\max_{q_2 \geq 0} \pi_2(q_1, q_2) = \max_{q_2 \geq 0} q_2[a - q_1 - q_2 - c],$$

which yields

$$R_2(q_1) = \frac{a - q_1 - c}{2},$$

provided $q_1 < a - c$. The same equation for $R_2(q_1)$ appeared in our analysis of the simultaneous-move Cournot game in Section 1.2.A. The difference is that here $R_2(q_1)$ is truly firm 2's reaction to firm 1's observed quantity, whereas in the Cournot analysis $R_2(q_1)$ is firm 2's best response to a hypothesized quantity to be simultaneously chosen by firm 1.

Since firm 1 can solve firm 2's problem as well as firm 2 can solve it, firm 1 should anticipate that the quantity choice q_1 will be met with the reaction $R_2(q_1)$. Thus, firm 1's problem in the first stage of the game amounts to

$$\begin{aligned} \max_{q_1 \geq 0} \pi_1(q_1, R_2(q_1)) &= \max_{q_1 \geq 0} q_1[a - q_1 - R_2(q_1) - c] \\ &= \max_{q_1 \geq 0} q_1 \frac{a - q_1 - c}{2}, \end{aligned}$$

which yields

$$q_1^* = \frac{a - c}{2} \quad \text{and} \quad R_2(q_1^*) = \frac{a - c}{4}$$

as the backwards-induction outcome of the Stackelberg duopoly game.[4]

---

[4] Just as "Cournot equilibrium" and "Bertrand equilibrium" typically refer to the Nash equilibria of the Cournot and Bertrand games, references to

Recall from Chapter 1 that in the Nash equilibrium of the Cournot game each firm produces $(a-c)/3$. Thus, aggregate quantity in the backwards-induction outcome of the Stackelberg game, $3(a-c)/4$, is greater than aggregate quantity in the Nash equilibrium of the Cournot game, $2(a-c)/3$, so the market-clearing price is lower in the Stackelberg game. In the Stackelberg game, however, firm 1 could have chosen its Cournot quantity, $(a-c)/3$, in which case firm 2 would have responded with its Cournot quantity. Thus, in the Stackelberg game, firm 1 could have achieved its Cournot profit level but chose to do otherwise, so firm 1's profit in the Stackelberg game must exceed its profit in the Cournot game. But the market-clearing price is lower in the Stackelberg game, so aggregate profits are lower, so the fact that firm 1 is better off implies that firm 2 is worse off in the Stackelberg than in the Cournot game.

The observation that firm 2 does worse in the Stackelberg than in the Cournot game illustrates an important difference between single- and multi-person decision problems. In single-person decision theory, having more information can never make the decision maker worse off. In game theory, however, having more information (or, more precisely, having it known to the other players that one has more information) *can* make a player worse off.

In the Stackelberg game, the information in question is firm 1's quantity: firm 2 knows q_1, and (as importantly) firm 1 knows that firm 2 knows q_1. To see the effect this information has, consider the modified sequential-move game in which firm 1 chooses q_1, after which firm 2 chooses q_2 but does so without observing q_1. If firm 2 believes that firm 1 has chosen its Stackelberg quantity $q_1^* = (a-c)/2$, then firm 2's best response is again $R_2(q_1^*) = (a-c)/4$. But if firm 1 anticipates that firm 2 will hold this belief and so choose this quantity, then firm 1 prefers to choose its best response to $(a-c)/4$—namely, $3(a-c)/8$—rather than its Stackelberg quantity $(a-c)/2$. Thus, firm 2 should not believe that firm 1 has chosen its Stackelberg quantity. Rather, the unique Nash equilibrium of this

---

"Stackelberg equilibrium" often mean that the game is sequential- rather than simultaneous-move. As noted in the previous section, however, sequential-move games sometimes have multiple Nash equilibria, only one of which is associated with the backwards-induction outcome of the game. Thus, "Stackelberg equilibrium" can refer both to the sequential-move nature of the game and to the use of a stronger solution concept than simply Nash equilibrium.

modified sequential-move game is for both firms to choose the quantity $(a-c)/3$—precisely the Nash equilibrium of the Cournot game, where the firms move simultaneously.[5] Thus, having firm 1 know that firm 2 knows q_1 hurts firm 2.

2.1.C Wages and Employment in a Unionized Firm

In Leontief's (1946) model of the relationship between a firm and a monopoly union (i.e., a union that is the monopoly seller of labor to the firm), the union has exclusive control over wages, but the firm has exclusive control over employment. (Similar qualitative conclusions emerge in a more realistic model in which the firm and the union bargain over wages but the firm retains exclusive control over employment.) The union's utility function is $U(w,L)$, where w is the wage the union demands from the firm and L is employment. Assume that $U(w,L)$ increases in both w and L. The firm's profit function is $\pi(w,L) = R(L) - wL$, where $R(L)$ is the revenue the firm can earn if it employs L workers (and makes the associated production and product-market decisions optimally). Assume that $R(L)$ is increasing and concave.

Suppose the timing of the game is: (1) the union makes a wage demand, w; (2) the firm observes (and accepts) w and then chooses employment, L; (3) payoffs are $U(w,L)$ and $\pi(w,L)$. We can say a great deal about the backwards-induction outcome of this game even though we have not assumed specific functional forms for $U(w,L)$ and $R(L)$ and so are not able to solve for this outcome explicitly.

First, we can characterize the firm's best response in stage (2), $L^*(w)$, to an arbitrary wage demand by the union in stage (1), w. Given w, the firm chooses $L^*(w)$ to solve

$$\max_{L\geq 0} \pi(w,L) = \max_{L\geq 0} R(L) - wL,$$

the first-order condition for which is

$$R'(L) - w = 0.$$

---

[5] This is an example of a claim we made in Section 1.1.A: in a normal-form game the players choose their strategies simultaneously, but this does not imply that the parties necessarily *act* simultaneously; it suffices that each choose his or her action without knowledge of the others' choices. For further discussion of this point, see Section 2.4.A.

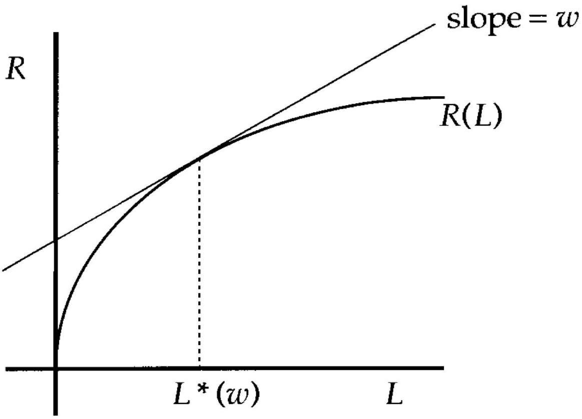

Figure 2.1.1.

To guarantee that the first-order condition $R'(L) - w = 0$ has a solution, assume that $R'(0) = \infty$ and that $R'(\infty) = 0$, as suggested in Figure 2.1.1.

Figure 2.1.2 plots $L^*(w)$ as a function of w (but uses axes that ease comparison with later figures) and illustrates that $L^*(w)$ cuts each of the firm's isoprofit curves at its maximum.[6] Holding L fixed, the firm does better when w is lower, so lower isoprofit curves represent higher profit levels. Figure 2.1.3 depicts the union's indifference curves. Holding L fixed, the union does better when w is higher, so higher indifference curves represent higher utility levels for the union.

We turn next to the union's problem at stage (1). Since the union can solve the firm's second-stage problem as well as the firm can solve it, the union should anticipate that the firm's reaction to the wage demand w will be to choose the employment level

---

[6] The latter property is merely a restatement of the fact that $L^*(w)$ maximizes $\pi(L, w)$ given w. If the union demands w', for example, then the firm's choice of L amounts to the choice of a point on the horizontal line $w = w'$. The highest feasible profit level is attained by choosing L such that the isoprofit curve through (L, w') is tangent to the constraint $w = w'$.

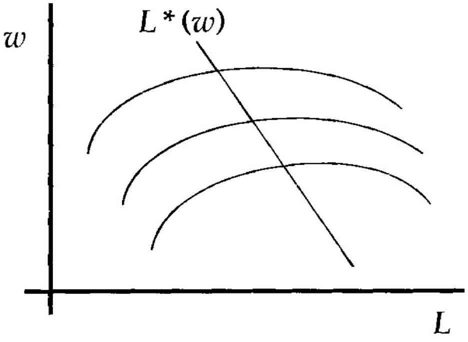

firm's isoprofit curves

Figure 2.1.2.

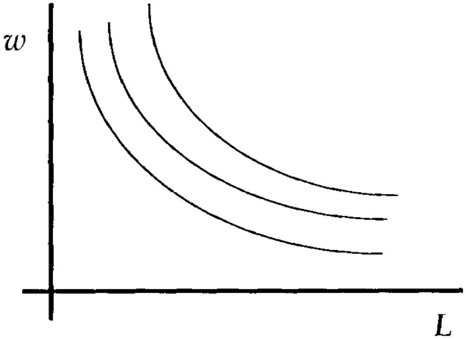

union's indifference curves

Figure 2.1.3.

$L^*(w)$. Thus, the union's problem at the first stage amounts to

$$\max_{w \geq 0} \; U(w, L^*(w)).$$

In terms of the indifference curves plotted in Figure 2.1.3, the

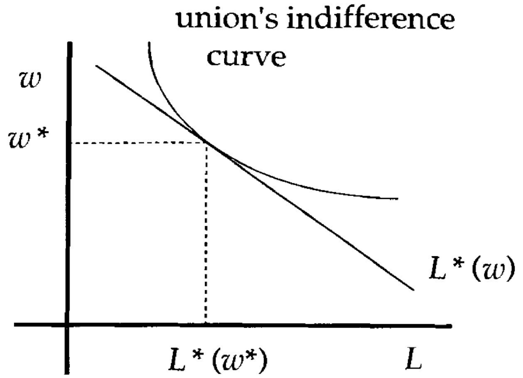

Figure 2.1.4.

union would like to choose the wage demand w that yields the outcome $(w, L^*(w))$ that is on the highest possible indifference curve. The solution to the union's problem is w^*, the wage demand such that the union's indifference curve through the point $(w^*, L^*(w^*))$ is tangent to $L^*(w)$ at that point; see Figure 2.1.4. Thus, $(w^*, L^*(w^*))$ is the backwards-induction outcome of this wage-and-employment game.

It is straightforward to see that $(w^*, L^*(w^*))$ is inefficient: both the union's utility and the firm's profit would be increased if w and L were in the shaded region in Figure 2.1.5. This inefficiency makes it puzzling that in practice firms seem to retain exclusive control over employment. (Allowing the firm and the union to bargain over the wage but leaving the firm with exclusive control over employment yields a similar inefficiency.) Espinosa and Rhee (1989) propose one answer to this puzzle, based on the fact that the union and the firm negotiate repeatedly over time (often every three years, in the United States). There may exist an equilibrium of such a repeated game in which the union's choice of w and the firm's choice of L lie in the shaded region of Figure 2.1.5, even though such values of w and L cannot arise as the backwards-induction outcome of a single negotiation. See Section 2.3 on repeated games and Problem 2.16 on the Espinosa-Rhee model.

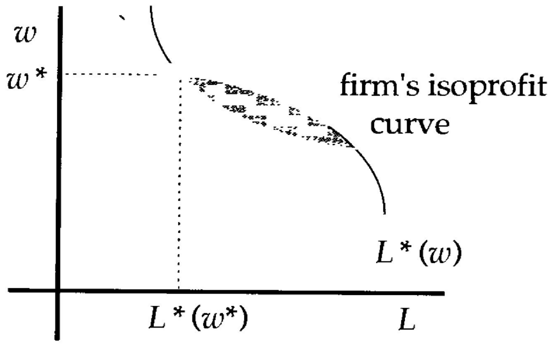

Figure 2.1.5.

2.1.D Sequential Bargaining

We begin with a three-period bargaining model from the class of games analyzed in Section 2.1.A. We then discuss Rubinstein's (1982) model, in which the number of periods is (potentially) infinite. In both models, settlement occurs immediately—protracted negotiations (such as strikes) do not occur. In Sobel and Takahashi's (1983) model of sequential bargaining under asymmetric information, in contrast, strikes occur with positive probability in the unique (perfect Bayesian) equilibrium; see Section 4.3.B.

Players 1 and 2 are bargaining over one dollar. They alternate in making offers: first player 1 makes a proposal that player 2 can accept or reject; if 2 rejects then 2 makes a proposal that 1 can accept or reject; and so on. Once an offer has been rejected, it ceases to be binding and is irrelevant to the subsequent play of the game. Each offer takes one period, and the players are impatient: they discount payoffs received in later periods by the factor δ per period, where $0 < \delta < 1$.[7]

---

[7] The *discount factor* δ reflects the time-value of money. A dollar received at the beginning of one period can be put in the bank to earn interest, say at rate r per period, and so will be worth $1+r$ dollars at the beginning of the next period. Equivalently, a dollar to be received at the beginning of the next period is worth only $1/(1+r)$ of a dollar now. Let $\delta = 1/(1+r)$. Then a payoff π to be received in the next period is worth only $\delta\pi$ now, a payoff π to be received two periods

A more detailed description of the timing of the three-period bargaining game is as follows.

(1a) At the beginning of the first period, player 1 proposes to take a share s_1 of the dollar, leaving $1 - s_1$ for player 2.

(1b) Player 2 either accepts the offer (in which case the game ends and the payoffs s_1 to player 1 and $1 - s_1$ to player 2 are immediately received) or rejects the offer (in which case play continues to the second period).

(2a) At the beginning of the second period, player 2 proposes that player 1 take a share s_2 of the dollar, leaving $1 - s_2$ for player 2. (Note the convention that s_t always goes to player 1, regardless of who made the offer.)

(2b) Player 1 either accepts the offer (in which case the game ends and the payoffs s_2 to player 1 and $1 - s_2$ to player 2 are immediately received) or rejects the offer (in which case play continues to the third period).

(3) At the beginning of the third period, player 1 receives a share s of the dollar, leaving $1 - s$ for player 2, where $0 < s < 1$.

In this three-period model, the third-period settlement $(s, 1 - s)$ is given exogenously. In the infinite-horizon model we later consider, the payoff s in the third period will represent player 1's payoff in the game that remains if the third period is reached (i.e., if the first two offers are rejected).

To solve for the backwards-induction outcome of this three-period game, we first compute player 2's optimal offer if the second period is reached. Player 1 can receive s in the third period by rejecting player 2's offer of s_2 this period, but the value this period of receiving s next period is only δs. Thus, player 1 will accept s_2 if and only if $s_2 \geq \delta s$. (We assume that each player will accept an offer if indifferent between accepting and rejecting.) Player 2's second-period decision problem therefore amounts to choosing

---

from now is worth only $\delta^2 \pi$ now, and so on. The value today of a future payoff is called the *present value* of that payoff.

between receiving $1-\delta s$ this period (by offering $s_2 = \delta s$ to player 1) and receiving $1 - s$ next period (by offering player 1 any $s_2 < \delta s$). The discounted value of the latter option is $\delta(1 - s)$, which is less than the $1 - \delta s$ available from the former option, so player 2's optimal second-period offer is $s_2^* = \delta s$. Thus, if play reaches the second period, player 2 will offer s_2^* and player 1 will accept.

Since player 1 can solve player 2's second-period problem as well as player 2 can, player 1 knows that player 2 can receive $1 - s_2^*$ in the second period by rejecting player 1's offer of s_1 this period, but the value this period of receiving $1 - s_2^*$ next period is only $\delta(1 - s_2^*)$. Thus, player 2 will accept $1 - s_1$ if and only if $1 - s_1 \geq \delta(1 - s_2^*)$, or $s_1 \leq 1 - \delta(1 - s_2^*)$. Player 1's first-period decision problem therefore amounts to choosing between receiving $1 - \delta(1 - s_2^*)$ this period (by offering $1 - s_1 = \delta(1 - s_2^*)$ to player 2) and receiving s_2^* next period (by offering any $1 - s_1 < \delta(1 - s_2^*)$ to player 2). The discounted value of the latter option is $\delta s_2^* = \delta^2 s$, which is less than the $1 - \delta(1 - s_2^*) = 1 - \delta(1 - \delta s)$ available from the former option, so player 1's optimal first-period offer is $s_1^* = 1 - \delta(1 - s_2^*) = 1 - \delta(1 - \delta s)$. Thus, in the backwards-induction outcome of this three-period game, player 1 offers the settlement $(s_1^*, 1 - s_1^*)$ to player 2, who accepts.

Now consider the infinite-horizon case. The timing is as described previously, except that the exogenous settlement in step (3) is replaced by an infinite sequence of steps (3a), (3b), (4a), (4b), and so on. Player 1 makes the offer in odd-numbered periods, player 2 in even-numbered; bargaining continues until one player accepts an offer. We would like to solve for the backwards-induction outcome of the infinite-horizon game by working backwards, as in all the applications analyzed so far. Because the game could go on infinitely, however, there is no last move at which to begin such an analysis. Fortunately, the following insight (first applied by Shaked and Sutton 1984) allows us to truncate the infinite-horizon game and apply the logic from the finite-horizon case: the game beginning in the third period (should it be reached) is identical to the game as a whole (beginning in the first period)—in both cases, player 1 makes the first offer, the players alternate in making subsequent offers, and the bargaining continues until one player accepts an offer.

Since we have not formally defined a backwards-induction outcome for this infinite-horizon bargaining game, our arguments will be informal (but can be made formal). Suppose that there is a backwards-induction outcome of the game as a whole in which players 1 and 2 receive the payoffs s and $1 - s$, respectively. We can use these payoffs in the game beginning in the third period, should it be reached, and then work backwards to the first period (as in the three-period model) to compute a new backwards-induction outcome for the game as a whole. In this new backwards-induction outcome, player 1 will offer the settlement $(f(s), 1 - f(s))$ in the first period and player 2 will accept, where $f(s) = 1 - \delta(1 - \delta s)$ is the share taken by player 1 in the first period of the three-period model above when the settlement $(s, 1 - s)$ is exogenously imposed in the third period.

Let s_H be the highest payoff player 1 can achieve in any backwards-induction outcome of the game as a whole. Imagine using s_H as the third-period payoff to player 1, as previously described: this will produce a new backwards-induction outcome in which player 1's first-period payoff is $f(s_H)$. Since $f(s) = 1 - \delta + \delta^2 s$ is increasing in s, $f(s_H)$ is the highest possible first-period payoff because s_H is the highest possible third-period payoff. But s_H is also the highest possible first-period payoff, so $f(s_H) = s_H$. Parallel arguments show that $f(s_L) = s_L$, where s_L is the lowest payoff player 1 can achieve in any backwards-induction outcome of the game as a whole. The only value of s that satisfies $f(s) = s$ is $1/(1+\delta)$, which we will denote by s^*. Thus, $s_H = s_L = s^*$, so there is a unique backwards-induction outcome in the game as a whole: in the first period, player 1 offers the settlement $(s^* = 1/(1 + \delta), 1 - s^* = \delta/(1 + \delta))$ to player 2, who accepts.

2.2 Two-Stage Games of Complete but Imperfect Information

2.2.A Theory: Subgame Perfection

We now enrich the class of games analyzed in the previous section. As in dynamic games of complete and perfect information, we continue to assume that play proceeds in a sequence of stages, with the moves in all previous stages observed before the next stage begins. Unlike in the games analyzed in the previous section, however, we now allow there to be simultaneous moves within each stage. As will be explained in Section 2.4, this simultaneity of moves within stages means that the games analyzed in this section have imperfect information. Nonetheless, these games share important features with the perfect-information games considered in the previous section.

We will analyze the following simple game, which we (uninspiredly!) call a two-stage game of complete but imperfect information:

1. Players 1 and 2 simultaneously choose actions a_1 and a_2 from feasible sets A_1 and A_2, respectively.

2. Players 3 and 4 observe the outcome of the first stage, (a_1, a_2), and then simultaneously choose actions a_3 and a_4 from feasible sets A_3 and A_4, respectively.

3. Payoffs are $u_i(a_1, a_2, a_3, a_4)$ for $i = 1, 2, 3, 4$.

Many economic problems fit this description.[8] Three examples (later discussed in detail) are bank runs, tariffs and imperfect international competition, and tournaments (e.g., competition among several vice presidents in a firm to be the next president). Other economic problems can be modeled by allowing for a longer sequence of stages, either by adding players or by allowing players to move in more than one stage. There could also be fewer players: in some applications, players 3 and 4 are players 1 and 2; in others, either player 2 or player 4 is missing.

We solve a game from this class by using an approach in the spirit of backwards induction, but this time the first step in working backwards from the end of the game involves solving a real game (the simultaneous-move game between players 3 and 4 in stage two, given the outcome from stage one) rather than solving a single-person optimization problem as in the previous section. To keep things simple, in this section we will assume that for each feasible outcome of the first-stage game, (a_1, a_2), the second-stage game that remains between players 3 and 4 has a unique Nash equilibrium, denoted by $(a_3^*(a_1, a_2), a_4^*(a_1, a_2))$. In Section 2.3.A (on repeated games) we consider the implications of relaxing this assumption.

---

[8] As in the previous section, the feasible action sets of players 3 and 4 in the second stage, A_3 and A_4, could be allowed to depend on the outcome of the first stage, (a_1, a_2). In particular, there may be values of (a_1, a_2) that end the game.

If players 1 and 2 anticipate that the second-stage behavior of players 3 and 4 will be given by $(a_3^*(a_1,a_2), a_4^*(a_1,a_2))$, then the first-stage interaction between players 1 and 2 amounts to the following simultaneous-move game:

1. Players 1 and 2 simultaneously choose actions a_1 and a_2 from feasible sets A_1 and A_2, respectively.

2. Payoffs are $u_i(a_1,a_2,a_3^*(a_1,a_2),a_4^*(a_1,a_2))$ for $i = 1,2$.

Suppose (a_1^*, a_2^*) is the unique Nash equilibrium of this simultaneous-move game. We will call $(a_1^*, a_2^*, a_3^*(a_1^*, a_2^*). a_4^*(a_1^*, a_2^*))$ the *subgame-perfect outcome* of this two-stage game. This outcome is the natural analog of the backwards-induction outcome in games of complete and perfect information, and the analogy applies to both the attractive and the unattractive features of the latter. Players 1 and 2 should not believe a threat by players 3 and 4 that the latter will respond with actions that are not a Nash equilibrium in the remaining second-stage game, because when play actually reaches the second stage at least one of players 3 and 4 will not want to carry out such a threat (exactly because it is not a Nash equilibrium of the game that remains at that stage). On the other hand, suppose that player 1 is also player 3, and that player 1 does not play a_1^* in the first stage: player 4 may then want to reconsider the assumption that player 3 (i.e., player 1) will play $a_3^*(a_1,a_2)$ in the second stage.

2.2.B Bank Runs

Two investors have each deposited D with a bank. The bank has invested these deposits in a long-term project. If the bank is forced to liquidate its investment before the project matures, a total of $2r$ can be recovered, where $D > r > D/2$. If the bank allows the investment to reach maturity, however, the project will pay out a total of $2R$, where $R > D$.

There are two dates at which the investors can make withdrawals from the bank: date 1 is before the bank's investment matures; date 2 is after. For simplicity, assume that there is no discounting. If both investors make withdrawals at date 1 then each receives r and the game ends. If only one investor makes a withdrawal at date 1 then that investor receives D, the other receives $2r - D$, and the game ends. Finally, if neither investor makes a withdrawal at date 1 then the project matures and the investors make withdrawal decisions at date 2. If both investors make withdrawals at date 2 then each receives R and the game ends. If only one investor makes a withdrawal at date 2 then that investor receives $2R - D$, the other receives D, and the game ends. Finally, if neither investor makes a withdrawal at date 2 then the bank returns R to each investor and the game ends.

In Section 2.4 we will discuss how to represent this game formally. For now, however, we will proceed informally. Let the payoffs to the two investors at dates 1 and 2 (as a function of their withdrawal decisions at these dates) be represented by the following pair of normal-form games. Note well that the normal-form game for date 1 is nonstandard: if both investors choose not to withdraw at date 1 then no payoff is specified; rather, the investors proceed to the normal-form game at date 2.

	withdraw	don't
withdraw	r, r	$D, 2r - D$
don't	$2r - D, D$	next stage

Date 1

	withdraw	don't
withdraw	R, R	$2R - D, D$
don't	$D, 2R - D$	R, R

Date 2

To analyze this game, we work backwards. Consider the normal-form game at date 2. Since $R > D$ (and so $2R - D > R$), "withdraw" strictly dominates "don't withdraw," so there is a unique Nash equilibrium in this game: both investors withdraw, leading to a payoff of (R, R). Since there is no discounting, we can simply substitute this payoff into the normal-form game at date 1, as in Figure 2.2.1. Since $r < D$ (and so $2r - D < r$), this one-period version of the two-period game has two pure-strategy Nash

	withdraw	don't
withdraw	r, r	$D, 2r - D$
don't	$2r - D, D$	R, R

Figure 2.2.1.

equilibria: (1) both investors withdraw, leading to a payoff of (r, r); (2) both investors do not withdraw, leading to a payoff of (R, R). Thus, the original two-period bank-runs game has two subgame-perfect outcomes (and so does not quite fit within the class of games defined in Section 2.2.A): (1) both investors withdraw at date 1, yielding payoffs of (r, r); (2) both investors do not withdraw at date 1 but do withdraw at date 2, yielding payoffs of (R, R) at date 2.

The first of these outcomes can be interpreted as a run on the bank. If investor 1 believes that investor 2 will withdraw at date 1 then investor 1's best response is to withdraw as well, even though both investors would be better off if they waited until date 2 to withdraw. This bank-run game differs from the Prisoners' Dilemma discussed in Chapter 1 in an important respect: both games have a Nash equilibrium that leads to a socially inefficient payoff; in the Prisoners' Dilemma this equilibrium is unique (and in dominant strategies), whereas here there also exists a second equilibrium that is efficient. Thus, this model does not predict when bank runs will occur, but does show that they can occur as an equilibrium phenomenon. See Diamond and Dybvig (1983) for a richer model.

2.2.C Tariffs and Imperfect International Competition

We turn next to an application from international economics. Consider two identical countries, denoted by $i = 1, 2$. Each country has a government that chooses a tariff rate, a firm that produces output for both home consumption and export, and consumers who buy on the home market from either the home firm or the foreign firm. If the total quantity on the market in country i is Q_i, then the market-clearing price is $P_i(Q_i) = a - Q_i$. The firm in country i (hereafter called firm i) produces h_i for home consumption and e_i for export. Thus, $Q_i = h_i + e_j$. The firms have a constant marginal cost, c, and no fixed costs. Thus, the total cost of production for firm i is $C_i(h_i, e_i) = c(h_i + e_i)$. The firms also incur tariff costs on exports: if firm i exports e_i to country j when government j has set the tariff rate t_j, then firm i must pay $t_j e_i$ to government j.

The timing of the game is as follows. First, the governments simultaneously choose tariff rates, t_1 and t_2. Second, the firms observe the tariff rates and simultaneously choose quantities for home consumption and for export, (h_1, e_1) and (h_2, e_2). Third, payoffs are profit to firm i and total welfare to government i, where total welfare to country i is the sum of the consumers' surplus[9] enjoyed by the consumers in country i, the profit earned by firm i, and the tariff revenue collected by government i from firm j:

$$\begin{aligned}\pi_i(t_i, t_j, h_i, e_i, h_j, e_j) &= [a - (h_i + e_j)]h_i + [a - (e_i + h_j)]e_i \\ &\quad - c(h_i + e_i) - t_j e_i, \\ W_i(t_i, t_j, h_i, e_i, h_j, e_j) &= \frac{1}{2}Q_i^2 + \pi_i(t_i, t_j, h_i, e_i, h_j, e_j) + t_i e_j.\end{aligned}$$

Suppose the governments have chosen the tariffs t_1 and t_2. If $(h_1^*, e_1^*, h_2^*, e_2^*)$ is a Nash equilibrium in the remaining (two-market) game between firms 1 and 2 then, for each i, (h_i^*, e_i^*) must solve

$$\max_{h_i, e_i \geq 0} \pi_i(t_i, t_j, h_i, e_i, h_j^*, e_j^*).$$

Since $\pi_i(t_i, t_j, h_i, e_i, h_j^*, e_j^*)$ can be written as the sum of firm i's profits on market i (which is a function of h_i and e_j^* alone) and firm i's profits on market j (which is a function of e_i, h_j^*, and t_j alone), firm i's two-market optimization problem simplifies into a pair of problems, one for each market: h_i^* must solve

$$\max_{h_i \geq 0} h_i[a - (h_i + e_j^*) - c],$$

and e_i^* must solve

$$\max_{e_i \geq 0} e_i[a - (e_i + h_j^*) - c] - t_j e_i.$$

---

[9] If a consumer buys a good for price p when she would have been willing to pay the value v, then she enjoys a surplus of $v - p$. Given the inverse demand curve $P_i(Q_i) = a - Q_i$, if the quantity sold on market i is Q_i, the aggregate consumer surplus can be shown to be $(1/2)Q_i^2$.

Assuming $e_j^* \leq a - c$, we have

$$h_i^* = \frac{1}{2}(a - e_j^* - c), \tag{2.2.1}$$

and assuming $h_j^* \leq a - c - t_j$, we have

$$e_i^* = \frac{1}{2}(a - h_j^* - c - t_j). \tag{2.2.2}$$

(The results we derive are consistent with both of these assumptions.) Both of the best-response functions (2.2.1) and (2.2.2) must hold for each $i = 1, 2$. Thus, we have four equations in the four unknowns $(h_1^*, e_1^*, h_2^*, e_2^*)$. Fortunately, these equations simplify into two sets of two equations in two unknowns. The solutions are

$$h_i^* = \frac{a - c + t_i}{3} \quad \text{and} \quad e_i^* = \frac{a - c - 2t_j}{3}. \tag{2.2.3}$$

Recall (from Section 1.2.A) that the equilibrium quantity chosen by both firms in the Cournot game is $(a - c)/3$, but that this result was derived under the assumption of symmetric marginal costs. In the equilibrium described by (2.2.3), in contrast, the governments' tariff choices make marginal costs asymmetric (as in Problem 1.6). On market i, for instance, firm i's marginal cost is c but firm j's is $c + t_i$. Since firm j's cost is higher it wants to produce less. But if firm j is going to produce less, then the market-clearing price will be higher, so firm i wants to produce more, in which case firm j wants to produce even less. Thus, in equilibrium, h_i^* increases in t_i and e_j^* decreases (at a faster rate) in t_i, as in (2.2.3).

Having solved the second-stage game that remains between the two firms after the governments choose tariff rates, we can now represent the first-stage interaction between the two governments as the following simultaneous-move game. First, the governments simultaneously choose tariff rates t_1 and t_2. Second, payoffs are $W_i(t_i, t_j, h_1^*, e_1^*, h_2^*, e_2^*)$ for government $i = 1, 2$, where h_i^* and e_i^* are functions of t_i and t_j as described in (2.2.3). We now solve for the Nash equilibrium of this game between the governments.

To simplify the notation, we will suppress the dependence of h_i^* on t_i and e_i^* on t_j: let $W_i^*(t_i, t_j)$ denote $W_i(t_i, t_j, h_1^*, e_1^*, h_2^*, e_2^*)$, the payoff to government i when it chooses the tariff rate t_i, government j chooses t_j, and firms i and j then play the Nash equilibrium given in (2.2.3). If (t_1^*, t_2^*) is a Nash equilibrium of this game between the governments then, for each i, t_i^* must solve

$$\max_{t_i \geq 0} W_i^*(t_i, t_j^*).$$

But $W_i^*(t_i, t_j^*)$ equals

$$\frac{(2(a-c)-t_i)^2}{18} + \frac{(a-c+t_i)^2}{9} + \frac{(a-c-2t_j^*)^2}{9} + \frac{t_i(a-c-2t_i)}{3},$$

so

$$t_i^* = \frac{a-c}{3}$$

for each i, independent of t_j^*. Thus, in this model, choosing a tariff rate of $(a-c)/3$ is a dominant strategy for each government. (In other models, such as when marginal costs are increasing, the governments' equilibrium strategies are not dominant strategies.) Substituting $t_i^* = t_j^* = (a-c)/3$ into (2.2.3) yields

$$h_i^* = \frac{4(a-c)}{9} \qquad \text{and} \qquad e_i^* = \frac{a-c}{9}$$

as the firms' quantity choices in the second stage. Thus, the subgame-perfect outcome of this tariff game is $(t_1^* = t_2^* = (a-c)/3$, $h_1^* = h_2^* = 4(a-c)/9$, $e_1^* = e_2^* = (a-c)/9)$.

In the subgame-perfect outcome the aggregate quantity on each market is $5(a-c)/9$. If the governments had chosen tariff rates equal to zero, however, then the aggregate quantity on each market would have been $2(a-c)/3$, just as in the Cournot model. Thus, the consumers' surplus on market i (which, as noted earlier, is simply one-half the square of the aggregate quantity on market i) is lower when the governments choose their dominant-strategy tariffs than it would be if they chose zero tariffs. In fact, zero tariffs are socially optimal, in the sense that $t_1 = t_2 = 0$ is the solution to

$$\max_{t_1, t_2 \geq 0} W_1^*(t_1, t_2) + W_2^*(t_2, t_1),$$

so there is an incentive for the governments to sign a treaty in which they commit to zero tariffs (i.e., free trade). (If negative tariffs—that is, subsidies—are feasible, the social optimum is for the governments to choose $t_1 = t_2 = -(a-c)$, which causes the home firm to produce zero for home consumption and to export the perfect-competition quantity to the other country.) Thus, given that firms i and j play the Nash equilibrium given in (2.2.3) in the second stage, the first-stage interaction between the governments is a Prisoners' Dilemma: the unique Nash equilibrium is in dominant strategies and is socially inefficient.

2.2.D Tournaments

Consider two workers and their boss. Worker i (where $i = 1$ or 2) produces output $y_i = e_i + \varepsilon_i$, where e_i is effort and ε_i is noise. Production proceeds as follows. First, the workers simultaneously choose nonnegative effort levels: $e_i \geq 0$. Second, the noise terms ε_1 and ε_2 are independently drawn from a density $f(\varepsilon)$ with zero mean. Third, the workers' outputs are observed but their effort choices are not. The workers' wages therefore can depend on their outputs but not (directly) on their efforts.

Suppose the workers' boss decides to induce effort from the workers by having them compete in a tournament, as first analyzed by Lazear and Rosen (1981).[10] The wage earned by the winner of the tournament (i.e., the worker with the higher output) is w_H; the wage earned by the loser is w_L. The payoff to a worker from earning wage w and expending effort e is $u(w, e) = w - g(e)$, where the disutility of effort, $g(e)$, is increasing and convex (i.e., $g'(e) > 0$ and $g''(e) > 0$). The payoff to the boss is $y_1 + y_2 - w_H - w_L$.

We now translate this application into the terms of the class of games discussed in Section 2.2.A. The boss is player 1, whose action a_1 is choosing the wages to be paid in the tournament, w_H and w_L. There is no player 2. The workers are players 3 and 4, who observe the wages chosen in the first stage and then simultaneously choose actions a_3 and a_4, namely the effort choices e_1 and e_2. (We later consider the possibility that, given the wages chosen by the boss, the workers prefer not to participate in the tournament and accept alternative employment instead.) Finally, the players' payoffs are as given earlier. Since outputs (and so also wages) are functions not only of the players actions but also of the noise terms ε_1 and ε_2, we work with the players' expected payoffs.

---

[10] To keep the exposition of this application simple, we ignore several technical details, such as conditions under which the worker's first-order condition is sufficient. Nonetheless, the analysis involves more probability than others thus far. The application can be skipped without loss of continuity.

Suppose the boss has chosen the wages w_H and w_L. If the effort pair (e_1^*, e_2^*) is to be a Nash equilibrium of the remaining game between the workers then, for each i, e_i^* must maximize worker i's expected wage, net of the disutility of effort: e_i^* must solve[11]

$$\begin{aligned} &\max_{e_i \geq 0} \; w_H \operatorname{Prob}\{y_i(e_i) > y_j(e_j^*)\} + w_L \operatorname{Prob}\{y_i(e_i) \leq y_j(e_j^*)\} - g(e_i) \\ &\qquad = (w_H - w_L) \operatorname{Prob}\{y_i(e_i) > y_j(e_j^*)\} + w_L - g(e_i), \end{aligned} \tag{2.2.4}$$

where $y_i(e_i) = e_i + \varepsilon_i$. The first-order condition for (2.2.4) is

$$(w_H - w_L) \frac{\partial \operatorname{Prob}\{y_i(e_i) > y_j(e_j^*)\}}{\partial e_i} = g'(e_i). \tag{2.2.5}$$

That is, worker i chooses e_i such that the marginal disutility of extra effort, $g'(e_i)$, equals the marginal gain from extra effort, which is the product of the wage gain from winning the tournament, $w_H - w_L$, and the marginal increase in the probability of winning.

By Bayes' rule,[12]

$$\begin{aligned} \operatorname{Prob}\{y_i(e_i) > y_j(e_j^*)\} &= \operatorname{Prob}\{\varepsilon_i > e_j^* + \varepsilon_j - e_i\} \\ &= \int_{\varepsilon_j} \operatorname{Prob}\{\varepsilon_i > e_j^* + \varepsilon_j - e_i \mid \varepsilon_j\} f(\varepsilon_j)\, d\varepsilon_j \\ &= \int_{\varepsilon_j} [1 - F(e_j^* - e_i + \varepsilon_j)] f(\varepsilon_j)\, d\varepsilon_j, \end{aligned}$$

---

[11] In writing (2.2.4), we assume that the noise density $f(\varepsilon)$ is such that the event that the workers' outputs are exactly equal happens with zero probability and so need not be considered in worker i's expected utility. (More formally, we assume that the density $f(\varepsilon)$ is atomless.) In a complete description of the tournament, it would be natural (but immaterial) to specify that the winner is determined by a coin flip or (equivalently, in this model) that both workers receive $(w_H + w_L)/2$.

[12] Bayes' rule provides a formula for $P(A \mid B)$, the (conditional) probability that an event A will occur given that an event B has already occurred. Let $P(A)$, $P(B)$, and $P(A, B)$ be the (prior) probabilities (i.e., the probabilities before either A or B has had a chance to take place) that A will occur, that B will occur, and that both A and B will occur, respectively. Bayes' rule states that $P(A \mid B) = P(A, B)/P(B)$. That is, the conditional probability of A given B equals the probability that both A and B will occur, divided by the prior probability that B will occur.

so the first-order condition (2.2.5) becomes

$$(w_H - w_L) \int_{\varepsilon_j} f(e_j^* - e_i + \varepsilon_j) f(\varepsilon_j)\, d\varepsilon_j = g'(e_i).$$

In a symmetric Nash equilibrium (i.e., $e_1^* = e_2^* = e^*$), we have

$$(w_H - w_L) \int_{\varepsilon_j} f(\varepsilon_j)^2\, d\varepsilon_j = g'(e^*). \tag{2.2.6}$$

Since $g(e)$ is convex, a bigger prize for winning (i.e., a larger value of $w_H - w_L$) induces more effort, as is intuitive. On the other hand, holding the prize constant, it is not worthwhile to work hard when output is very noisy, because the outcome of the tournament is likely to be determined by luck rather than effort. If ε is normally distributed with variance σ^2, for example, then

$$\int_{\varepsilon_j} f(\varepsilon_j)^2\, d\varepsilon_j = \frac{1}{2\sigma\sqrt{\pi}},$$

which decreases in σ, so e^* indeed decreases in σ.

We now work backwards to the first stage of the game. Suppose that if the workers agree to participate in the tournament (rather than accept alternative employment) then they will respond to the wages w_H and w_L by playing the symmetric Nash equilibrium characterized by (2.2.6). (We thus ignore the possibilities of asymmetric equilibria and of an equilibrium in which the workers' effort choices are given by the corner solution $e_1 = e_2 = 0$, rather than by the first-order condition (2.2.5).) Suppose also that the workers' alternative employment opportunity would provide utility U_a. Since in the symmetric Nash equilibrium each worker wins the tournament with probability one-half (i.e., $\text{Prob}\{y_i(e^*) > y_j(e^*)\} = 1/2$), if the boss intends to induce the workers to participate in the tournament then she must choose wages that satisfy

$$\frac{1}{2}w_H + \frac{1}{2}w_L - g(e^*) \geq U_a. \tag{2.2.7}$$

Assuming that U_a is low enough that the boss wants to induce the workers to participate in the tournament, she therefore chooses wages to maximize expected profit, $2e^* - w_H - w_L$, subject to (2.2.7). At the optimum, (2.2.7) holds with equality:

$$w_L = 2U_a + 2g(e^*) - w_H. \tag{2.2.8}$$

Expected profit then becomes $2e^* - 2U_a - 2g(e^*)$, so the boss wishes to choose wages such that the induced effort, e^*, maximizes $e^* - g(e^*)$. The optimal induced effort therefore satisfies the first-order condition $g'(e^*) = 1$. Substituting this into (2.2.6) implies that the optimal prize, $w_H - w_L$, solves

$$(w_H - w_L) \int_{\varepsilon_j} f(\varepsilon_j)^2 \, d\varepsilon_j = 1,$$

and (2.2.8) then determines w_H and w_L themselves.

2.3 Repeated Games

In this section we analyze whether threats and promises about future behavior can influence current behavior in repeated relationships. Much of the intuition is given in the two-period case; a few ideas require an infinite horizon. We also define subgame-perfect Nash equilibrium for repeated games. This definition is simpler to express for the special case of repeated games than for the general dynamic games of complete information we consider in Section 2.4.B. We introduce it here so as to ease the exposition later.

2.3.A Theory: Two-Stage Repeated Games

Consider the Prisoners' Dilemma given in normal form in Figure 2.3.1. Suppose two players play this simultaneous-move game twice, observing the outcome of the first play before the second play begins, and suppose the payoff for the entire game is simply the sum of the payoffs from the two stages (i.e., there is no

		Player 2	
		L_2	R_2
Player 1	L_1	1,1	5,0
	R_1	0,5	4,4

Figure 2.3.1.

		Player 2	
		L_2	R_2
Player 1	L_1	2,2	6,1
	R_1	1,6	5,5

Figure 2.3.2.

discounting). We will call this repeated game the two-stage Prisoners' Dilemma. It belongs to the class of games analyzed in Section 2.2.A. Here players 3 and 4 are identical to players 1 and 2, the action spaces A_3 and A_4 are identical to A_1 and A_2, and the payoffs $u_i(a_1, a_2, a_3, a_4)$ are simply the sum of the payoff from the first-stage outcome (a_1, a_2) and the payoff from the second-stage outcome (a_3, a_4). Furthermore, the two-stage Prisoners' Dilemma satisfies the assumption we made in Section 2.2.A: for each feasible outcome of the first-stage game, (a_1, a_2), the second-stage game that remains between players 3 and 4 has a unique Nash equilibrium, denoted by $(a_3^*(a_1, a_2), a_4^*(a_1, a_2))$. In fact, the two-stage Prisoners' Dilemma satisfies this assumption in the following stark way. In Section 2.2.A we allowed for the possibility that the Nash equilibrium of the remaining second-stage game depends on the first-stage outcome—hence the notation $(a_3^*(a_1, a_2), a_4^*(a_1, a_2))$ rather than simply (a_3^*, a_4^*). (In the tariff game, for example, the firms' equilibrium quantity choices in the second stage depend on the governments' tariff choices in the first stage.) In the two-stage Prisoners' Dilemma, however, the unique equilibrium of the second-stage game is (L_1, L_2), regardless of the first-stage outcome.

Following the procedure described in Section 2.2.A for computing the subgame-perfect outcome of such a game, we analyze the first stage of the two-stage Prisoners' Dilemma by taking into account that the outcome of the game remaining in the second stage will be the Nash equilibrium of that remaining game—namely, (L_1, L_2) with payoff $(1, 1)$. Thus, the players' first-stage interaction in the two-stage Prisoners' Dilemma amounts to the one-shot game in Figure 2.3.2, in which the payoff pair $(1, 1)$ for the second stage has been added to each first-stage payoff pair. The game in Figure 2.3.2 also has a unique Nash equilibrium: (L_1, L_2). Thus, the unique subgame-perfect outcome of the two-stage Prisoners'

Dilemma is (L_1, L_2) in the first stage, followed by (L_1, L_2) in the second stage. Cooperation—that is, (R_1, R_2)—cannot be achieved in either stage of the subgame-perfect outcome.

This argument holds more generally. (Here we temporarily depart from the two-period case to allow for any finite number of repetitions, T.) Let $G = \{A_1, \ldots, A_n; u_1, \ldots, u_n\}$ denote a static game of complete information in which players 1 through n simultaneously choose actions a_1 through a_n from the action spaces A_1 through A_n, respectively, and payoffs are $u_1(a_1, \ldots, a_n)$ through $u_n(a_1, \ldots, a_n)$. The game G will be called the *stage game* of the repeated game.

Definition *Given a stage game G, let G(T) denote the* ***finitely repeated game*** *in which G is played T times, with the outcomes of all preceding plays observed before the next play begins. The payoffs for G(T) are simply the sum of the payoffs from the T stage games.*

Proposition *If the stage game G has a unique Nash equilibrium then, for any finite T, the repeated game G(T) has a unique subgame-perfect outcome: the Nash equilibrium of G is played in every stage.*[13]

We now return to the two-period case, but consider the possibility that the stage game G has multiple Nash equilibria, as in Figure 2.3.3. The strategies labeled L_i and M_i mimic the Prisoners' Dilemma from Figure 2.3.1, but the strategies labeled R_i have been added to the game so that there are now two pure-strategy Nash equilibria: (L_1, L_2), as in the Prisoners' Dilemma, and now also (R_1, R_2). It is of course artificial to add an equilibrium to the Prisoners' Dilemma in this way, but our interest in this game is expositional rather than economic. In the next section we will see that infinitely repeated games share this multiple-equilibria spirit even if the stage game being repeated infinitely has a unique Nash equilibrium, as does the Prisoners' Dilemma. Thus, in this section we

---

[13] Analogous results hold if the stage game G is a dynamic game of complete information. Suppose G is a dynamic game of complete and perfect information from the class defined in Section 2.1.A. If G has a unique backwards-induction outcome, then $G(T)$ has a unique subgame-perfect outcome: the backwards-induction outcome of G is played in every stage. Similarly, suppose G is a two-stage game from the class defined in Section 2.2.A. If G has a unique subgame-perfect outcome, then $G(T)$ has a unique subgame-perfect outcome: the subgame-perfect outcome of G is played in every stage.

	L_2	M_2	R_2
L_1	1,1	5,0	0,0
M_1	0,5	4,4	0,0
R_1	0,0	0,0	3,3

Figure 2.3.3.

analyze an artificial stage game in the simple two-period framework, and thereby prepare for our later analysis of an economically interesting stage game in the infinite-horizon framework.

Suppose the stage game in Figure 2.3.3 is played twice, with the first-stage outcome observed before the second stage begins. We will show that there is a subgame-perfect outcome of this repeated game in which the strategy pair (M_1, M_2) is played in the first stage.[14] As in Section 2.2.A, assume that in the first stage the players anticipate that the second-stage outcome will be a Nash equilibrium of the stage game. Since this stage game has more than one Nash equilibrium, it is now possible for the players to anticipate that different first-stage outcomes will be followed by different stage-game equilibria in the second stage. Suppose, for example, that the players anticipate that (R_1, R_2) will be the second-stage outcome if the first-stage outcome is (M_1, M_2), but that (L_1, L_2) will be the second-stage outcome if any of the eight other first-stage outcomes occurs. The players' first-stage interaction then amounts to the one-shot game in Figure 2.3.4, where $(3,3)$ has been added to the (M_1, M_2)-cell and $(1,1)$ has been added to the eight other cells.

There are three pure-strategy Nash equilibria in the game in Figure 2.3.4: (L_1, L_2), (M_1, M_2), and (R_1, R_2). As in Figure 2.3.2,

---

[14] Strictly speaking, we have defined the notion of a subgame-perfect outcome only for the class of games defined in Section 2.2.A. The two-stage Prisoner's Dilemma belongs to this class because for each feasible outcome of the first-stage game there is a unique Nash equilibrium of the remaining second-stage game. The two-stage repeated game based on the stage game in Figure 2.3.3 does not belong to this class, however, because the stage game has multiple Nash equilibria. We will not formally extend the definition of a subgame-perfect outcome so that it applies to all two-stage repeated games, both because the change in the definition is minuscule and because even more general definitions appear in Sections 2.3.B and 2.4.B.

	L_2	M_2	R_2
L_1	2,2	6,1	1,1
M_1	1,6	7,7	1,1
R_1	1,1	1,1	4,4

Figure 2.3.4.

Nash equilibria of this one-shot game correspond to subgame-perfect outcomes of the original repeated game. Let $((w,x),(y,z))$ denote an outcome of the repeated game—(w,x) in the first stage and (y,z) in the second. The Nash equilibrium (L_1, L_2) in Figure 2.3.4 corresponds to the subgame-perfect outcome $((L_1,L_2),(L_1,L_2))$ in the repeated game, because the anticipated second-stage outcome is (L_1, L_2) following anything but (M_1, M_2) in the first stage. Likewise, the Nash equilibrium (R_1, R_2) in Figure 2.3.4 corresponds to the subgame-perfect outcome $((R_1,R_2),(L_1,L_2))$ in the repeated game. These two subgame-perfect outcomes of the repeated game simply concatenate Nash equilibrium outcomes from the stage game, but the third Nash equilibrium in Figure 2.3.4 yields a qualitatively different result: (M_1, M_2) in Figure 2.3.4 corresponds to the subgame-perfect outcome $((M_1,M_2),(R_1,R_2))$ in the repeated game, because the anticipated second-stage outcome is (R_1, R_2) following (M_1, M_2). Thus, as claimed earlier, cooperation can be achieved in the first stage of a subgame-perfect outcome of the repeated game. This is an example of a more general point: if $G = \{A_1,\ldots,A_n;u_1,\ldots,u_n\}$ is a static game of complete information with multiple Nash equilibria then there may be subgame-perfect outcomes of the repeated game $G(T)$ in which, for any $t < T$, the outcome in stage t is not a Nash equilibrium of G. We return to this idea in the infinite-horizon analysis in the next section.

The main point to extract from this example is that credible threats or promises about future behavior can influence current behavior. A second point, however, is that subgame-perfection may not embody a strong enough definition of credibility. In deriving the subgame-perfect outcome $((M_1,M_2),(R_1,R_2))$, for example, we assumed that the players anticipate that (R_1, R_2) will be the second-stage outcome if the first-stage outcome is (M_1, M_2)

	L_2	M_2	R_2	P_2	Q_2
L_1	1,1	5,0	0,0	0,0	0,0
M_1	0,5	4,4	0,0	0,0	0,0
R_1	0,0	0,0	3,3	0,0	0,0
P_1	0,0	0,0	0,0	$4,\frac{1}{2}$	0,0
Q_1	0,0	0,0	0,0	0,0	$\frac{1}{2},4$

Figure 2.3.5.

and that (L_1, L_2) will be the second-stage outcome if any of the eight other first-stage outcomes occurs. But playing (L_1, L_2) in the second stage, with its payoff of $(1,1)$, may seem silly when (R_1, R_2), with its payoff of $(3,3)$, is also available as a Nash equilibrium of the remaining stage game. Loosely put, it would seem natural for the players to renegotiate.[15] If (M_1, M_2) does not occur as the first-stage outcome, so that (L_1, L_2) is supposed to be played in the second stage, then each player might reason that bygones are bygones and that the unanimously preferred stage-game equilibrium (R_1, R_2) should be played instead. But if (R_1, R_2) is to be the second-stage outcome after every first-stage outcome, then the incentive to play (M_1, M_2) in the first stage is destroyed: the first-stage interaction between the two players simply amounts to the one-shot game in which the payoff $(3,3)$ has been added to each cell of the stage game in Figure 2.3.3, so L_i is player i's best response to M_j.

To suggest a solution to this renegotiation problem, we consider the game in Figure 2.3.5, which is even more artificial than the game in Figure 2.3.3. Once again, our interest in this game is expositional rather than economic. The ideas we develop here to address renegotiation in this artificial game can also be applied to renegotiation in infinitely repeated games; see Farrell and Maskin (1989), for example.

---

[15] This is loose usage because "renegotiate" suggests that communication (or even bargaining) occurs between the first and second stages. If such actions are possible, then they should be included in the description and analysis of the game. Here we assume that no such actions are possible, so by "renegotiate" we have in mind an analysis based on introspection.

This stage game adds the strategies P_i and Q_i to the stage game in Figure 2.3.3. There are four pure-strategy Nash equilibria of the stage game: (L_1, L_2) and (R_1, R_2), and now also (P_1, P_2) and (Q_1, Q_2). As before, the players unanimously prefer (R_1, R_2) to (L_1, L_2). More importantly, there is no Nash equilibrium (x, y) in Figure 2.3.5 such that the players unanimously prefer (x, y) to (P_1, P_2), or (Q_1, Q_2), or (R_1, R_2). We say that (R_1, R_2) *Pareto-dominates* (L_1, L_2), and that (P_1, P_2), (Q_1, Q_2), and (R_1, R_2) are on the *Pareto frontier* of the payoffs to Nash equilibria of the stage game in Figure 2.3.5.

Suppose the stage game in Figure 2.3.5 is played twice, with the first-stage outcome observed before the second stage begins. Suppose further that the players anticipate that the second-stage outcome will be as follows: (R_1, R_2) if the first-stage outcome is (M_1, M_2); (P_1, P_2) if the first-stage outcome is (M_1, w), where w is anything but M_2; (Q_1, Q_2) if the first-stage outcome is (x, M_2), where x is anything but M_1; and (R_1, R_2) if the first-stage outcome is (y, z), where y is anything but M_1 and z is anything but M_2. Then $((M_1, M_2), (R_1, R_2))$ is a subgame-perfect outcome of the repeated game, because each player gets $4 + 3$ from playing M_i and then R_i but only $5 + 1/2$ from deviating to L_i in the first stage (and even less from other deviations). More importantly, the difficulty in the previous example does not arise here. In the two-stage repeated game based on Figure 2.3.3, the only way to punish a player for deviating in the first stage was to play a Pareto-dominated equilibrium in the second stage, thereby also punishing the punisher. Here, in contrast, there are three equilibria on the Pareto frontier—one to reward good behavior by both players in the first stage, and two others to be used not only to punish a player who deviates in the first stage but also to reward the punisher. Thus, if punishment is called for in the second stage, there is no other stage-game equilibrium the punisher would prefer, so the punisher cannot be persuaded to renegotiate the punishment.

2.3.B Theory: Infinitely Repeated Games

We now turn to infinitely repeated games. As in the finite-horizon case, the main theme is that credible threats or promises about future behavior can influence current behavior. In the finite-horizon case we saw that if there are multiple Nash equilibria of the stage game G then there may be subgame-perfect outcomes of the repeated game $G(T)$ in which, for any $t < T$, the outcome of stage t is not a Nash equilibrium of G. A stronger result is true in infinitely repeated games: even if the stage game has a unique Nash equilibrium, there may be subgame-perfect outcomes of the infinitely repeated game in which no stage's outcome is a Nash equilibrium of G.

We begin by studying the infinitely repeated Prisoners' Dilemma. We then consider the class of infinitely repeated games analogous to the class of finitely repeated games defined in the previous section: a static game of complete information, G, is repeated infinitely, with the outcomes of all previous stages observed before the current stage begins. For these classes of finitely and infinitely repeated games, we define a player's strategy, a subgame, and a subgame-perfect Nash equilibrium. (In Section 2.4.B we define these concepts for general dynamic games of complete information, not just for these classes of repeated games.) We then use these definitions to state and prove Friedman's (1971) Theorem (also called the Folk Theorem).[16]

Suppose the Prisoners' Dilemma in Figure 2.3.6 is to be repeated infinitely and that, for each t, the outcomes of the $t-1$ preceding plays of the stage game are observed before the t^{th} stage begins. Simply summing the payoffs from this infinite sequence of stage games does not provide a useful measure of a player's payoff in the infinitely repeated game. Receiving a payoff of 4 in every period is better than receiving a payoff of 1 in every period, for example, but the sum of the payoffs is infinity in both cases. Recall (from Rubinstein's bargaining model in Section 2.1.D) that the discount factor $\delta = 1/(1+r)$ is the value today of a dollar to be received one stage later, where r is the interest rate per stage. Given a discount factor and a player's payoffs from an infinite

---

[16] The original Folk Theorem concerned the payoffs of all the Nash equilibria of an infinitely repeated game. This result was called the Folk Theorem because it was widely known among game theorists in the 1950s, even though no one had published it. Friedman's (1971) Theorem concerns the payoffs of certain subgame-perfect Nash equilibria of an infinitely repeated game, and so strengthens the original Folk Theorem by using a stronger equilibrium concept—subgame-perfect Nash equilibrium rather than Nash equilibrium. The earlier name has stuck, however: Friedman's Theorem (and later results) are sometimes called Folk Theorems, even though they were not widely known among game theorists before they were published.

Player 1 \ Player 2	L_2	R_2
L_1	1,1	5,0
R_1	0,5	4,4

Figure 2.3.6.

sequence of stage games, we can compute the *present value* of the payoffs—the lump-sum payoff that could be put in the bank now so as to yield the same bank balance at the end of the sequence.

Definition *Given the discount factor δ, the* ***present value*** *of the infinite sequence of payoffs $\pi_1, \pi_2, \pi_3, \ldots$ is*

$$\pi_1 + \delta\pi_2 + \delta^2\pi_3 + \cdots = \sum_{t=1}^{\infty} \delta^{t-1}\pi_t.$$

We can also use δ to reinterpret what we call an infinitely repeated game as a repeated game that ends after a random number of repetitions. Suppose that after each stage is played a (weighted) coin is flipped to determine whether the game will end. If the probability is p that the game ends immediately, and therefore $1-p$ that the game continues for at least one more stage, then a payoff π to be received in the next stage (if it is played) is worth only $(1-p)\pi/(1+r)$ before this stage's coin flip occurs. Likewise, a payoff π to be received two stages from now (if both it and the intervening stage are played) is worth only $(1-p)^2\pi/(1+r)^2$ before this stage's coin flip occurs. Let $\delta = (1-p)/(1+r)$. Then the present value $\pi_1 + \delta\pi_2 + \delta^2\pi_3 + \cdots$ reflects both the time-value of money and the possibility that the game will end.

Consider the infinitely repeated Prisoners' Dilemma in which each player's discount factor is δ and each player's payoff in the repeated game is the present value of the player's payoffs from the stage games. We will show that cooperation—that is, (R_1, R_2)—can occur in every stage of a subgame-perfect outcome of the infinitely repeated game, even though the only Nash equilibrium in the stage game is noncooperation—that is, (L_1, L_2). The argument is in the spirit of our analysis of the two-stage repeated game based on Figure 2.3.3 (the stage game in which we added a second Nash equilibrium to the Prisoners' Dilemma): if the players cooperate today then they play a high-payoff equilibrium tomorrow; otherwise they play a low-payoff equilibrium tomorrow. The difference between the two-stage repeated game and the infinitely repeated game is that here the high-payoff equilibrium that might be played tomorrow is not artificially added to the stage game but rather represents continuing to cooperate tomorrow and thereafter.

Suppose player i begins the infinitely repeated game by cooperating and then cooperates in each subsequent stage game if and only if both players have cooperated in every previous stage. Formally, player i's strategy is:

> Play R_i in the first stage. In the t^{th} stage, if the outcome of all $t-1$ preceding stages has been (R_1, R_2) then play R_i; otherwise, play L_i.

This strategy is an example of a *trigger strategy*, so called because player i cooperates until someone fails to cooperate, which triggers a switch to noncooperation forever after. If both players adopt this trigger strategy then the outcome of the infinitely repeated game will be (R_1, R_2) in every stage. We first argue that if δ is close enough to one then it is a Nash equilibrium of the infinitely repeated game for both players to adopt this strategy. We then argue that such a Nash equilibrium is subgame-perfect, in a sense to be made precise.

To show that it is a Nash equilibrium of the infinitely repeated game for both players to adopt the trigger strategy, we will assume that player i has adopted the trigger strategy and then show that, provided δ is close enough to one, it is a best response for player j to adopt the strategy also. Since player i will play L_i forever once one stage's outcome differs from (R_1, R_2), player j's best response is indeed to play L_j forever once one stage's outcome differs from (R_1, R_2). It remains to determine player j's best response in the first stage, and in any stage such that all the preceding outcomes have been (R_1, R_2). Playing L_j will yield a payoff of 5 this stage but will trigger noncooperation by player i (and therefore also by player j) forever after, so the payoff in every future stage will be 1. Since $1+\delta+\delta^2+\cdots = 1/(1-\delta)$, the present value of this sequence of payoffs is

$$5+\delta\cdot 1+\delta^2\cdot 1+\cdots = 5+\frac{\delta}{1-\delta}\,.$$

Alternatively, playing R_j will yield a payoff of 4 in this stage and will lead to exactly the same choice between L_j and R_j in the next stage. Let V denote the present value of the infinite sequence of payoffs player j receives from making this choice optimally (now and every time it arises subsequently). If playing R_j is optimal then

$$V = 4 + \delta V,$$

or $V = 4/(1-\delta)$, because playing R_j leads to the same choice next stage. If playing L_j is optimal then

$$V = 5 + \frac{\delta}{1-\delta},$$

as derived earlier. So playing R_j is optimal if and only if

$$\frac{4}{1-\delta} \geq 5 + \frac{\delta}{1-\delta}, \tag{2.3.1}$$

or $\delta \geq 1/4$. Thus, in the first stage, and in any stage such that all the preceding outcomes have been (R_1, R_2), player j's optimal action (given that player i has adopted the trigger strategy) is R_j if and only if $\delta \geq 1/4$. Combining this observation with the fact that j's best response is to play L_j forever once one stage's outcome differs from (R_1, R_2), we have that it is a Nash equilibrium for both players to play the trigger strategy if and only if $\delta \geq 1/4$.

We now want to argue that such a Nash equilibrium is subgame-perfect. To do so, we define a strategy in a repeated game, a subgame in a repeated game, and a subgame-perfect Nash equilibrium in a repeated game. In order to illustrate these concepts with simple examples from the previous section, we will define them for both finitely and infinitely repeated games. In the previous section we defined the finitely repeated game $G(T)$ based on a stage game $G = \{A_1, \ldots, A_n; u_1, \ldots, u_n\}$—a static game of complete information in which players 1 through n simultaneously choose actions a_1 through a_n from the action spaces A_1 through A_n, respectively, and payoffs are $u_1(a_1, \ldots, a_n)$ through $u_n(a_1, \ldots, a_n)$. We now define the analogous infinitely repeated game.[17]

---

[17] One can of course also define a repeated game based on a dynamic stage game. In this section we restrict attention to static stage games so as to present the main ideas in a simple way. The applications in Sections 2.3.D and 2.3.E are repeated games based on dynamic stage games.